Reactor Antineutrino Oscillation Detection in Liquid Scintillator at SNO+

Using (α, n)-IBD Pulse Shape Discrimination and Matter Effect Oscillation to Improve the Measurement of Long Baseline Neutrino Oscillation

James Page

A thesis presented for the degree of
Doctor of Philosophy

UNIVERSITY
OF SUSSEX

Department of Mathematical and Physical Sciences
University of Sussex
United Kingdom
September 2024

Reactor Antineutrino Oscillation Measurement at SNO+ in Liquid Scintillator

Using (α, n)-IBD Pulse Shape Discrimination and Matter Effect Oscillation to Improve the Measurement of Long Baseline Neutrino Oscillation

James Page

Abstract

The SNO+ experiment is a large liquid scintillator detector deep underground, with the main goal of measuring the speculative neutrino-less double-beta decay. It is the successor to the original Sudbury Neutrino Observatory (SNO), and, just like its predecessor, is also used to study long baseline neutrino oscillation. However, unlike SNO, it can do so by measuring inverse beta decays (IBD) induced by electron antineutrinos from nearby nuclear reactors. Detection of these has already been demonstrated when SNO+ was filled with ultra-pure water, and when half of this was replaced with liquid scintillator. The analysis presented here uses the currently largest available dataset from SNO+, harnessing 134.5 days of livetime from a detector fully loaded with liquid scintillator. A detailed review of the predictions of the IBD flux and its backgrounds is performed, teasing out any sources of uncertainty and constraining them where possible. The signal is then purified from these backgrounds with various cuts and coincidence tagging, followed by a new pulse shape discrimination classifier based on the Fisher discriminant. In simulations, this classifier achieves a selection efficiency of the primary background – (α, n) induced proton recoils – of 11.2%, while only sacrificing 3.9% of the total reactor IBD signal. The coincidence tagging produces 57 candidate events from the 134.5 days of data, of which 13 are removed by the classifier in a way consistent with predictions. An oscillation analysis is performed on both these datasets, using an efficient fitting framework allowing both shape and normalisation systematics to be floated alongside the oscillation parameters. This oscillation is realised via a novel algorithm which accounts for the full effect of three-flavour neutrino oscillation while propagating through matter of a constant density. The total expected flux is corrected by 0.6% compared to the vacuum case, while achieving competitive computational efficiency. A measurement of $\Delta m^2_{21} = \left(8.06^{+0.50}_{-0.44}\right) \times 10^{-5}\text{eV}^2$ is reported, making SNO+ hold the world's second most precise measurement behind KamLAND. Meanwhile, the measurement of $\sin^2\theta_{12} = 0.43^{+0.19}_{-0.33}$ is found to be in agreement with previous results. The geo-neutrino flux at SNO+ was simultaneously fit to 70^{+44}_{-40} TNU, making it tentatively the third location such a measurement has been performed worldwide. The (α, n)-IBD classifier is found to not improve the oscillation analysis at this stage due to the limited sample size. However, a sensitivity analysis demonstrates that it is expected to play a critical role in producing a world leading measurement of Δm^2_{21} in the coming one to two years, while enabling far more accurate measurements of the geo-neutrino flux as data accumulates.

“It’s a dangerous business, Frodo, going out your door. You step onto the road, and if you don’t keep your feet, there’s no knowing where you might be swept off to.”
Bilbo Baggins – *The Lord of the Rings*, J.R.R. Tolkien

To my cat.

Declaration

The author declares that this thesis, whether in the same or a different form, has not previously been submitted to this or any other university for a degree. The work presented here was carried out as part of the SNO+ collaboration, from whose members much information and many practices were derived. Whether work was performed by the author or not is generally indicated throughout the text. However, in broad strokes, chapters 2 and 3 contain no original work, while 4 is a mix of a pre-existing framework and extra contributions. In particular, the systematic uncertainties are reviewed in detail for the first time, as well as new antineutrino spectra introduced. Chapter 5 is mostly a literature review, but also contains a new oscillation algorithm developed solely by the author, which was already published during the thesis. While chapter 6 is mostly based on previous analyses, all the work was performed by the author except where otherwise stated. Finally, chapters 7 and 8 are almost entirely original work. However, even the entirely original work generally takes inspiration from previous iterations of similar analyses by others at SNO+ as its starting point (in particular works by Refs. [1, 2, 3]), and benefited from sound council.

Acknowledgements

I would like to thank everyone at SNO+ for making it a welcoming and friendly environment. The young member socials and wonderful on-site team made the Covid years more bearable, and my LTA a wonderful experience. Rafael, Daniel and Katharine made for great flat-mates and friends all the way out in Sudbury, and produced a truly unforgettable summer. I would also like to thank my supervisor Lisa for seeing my PhD application and giving me a chance, even with my atypical profile. She has been the most encouraging, kind and supportive supervisor I could ask for, with a surprising amount in common. Charlie and Sammy helped introduce me to the University of Sussex and SNO+, and provided invaluable guidance in my first two years. Thank you both for bearing with me as I pestered you mercilessly. The Bubble Chamber Cup and the most Canadian road trip to Kingston will never be forgotten! I must also thank my family for supporting and encouraging me through this whole journey, taking an interest in my work and always being proud of me. Lastly, thank you Victoria for spending that unforgettable summer with me, and encouraging and supporting me through the main part of my PhD. Thank you for putting up with my endless work every time I needed to “just finish something” and nothing was working, and still being excited about the silly things in life. I could not have made it this far without all of these people making it worthwhile. Speaking of endless work, thank you to the last minute debugging assist from Christopher, and rapid proof-reading from Ana-Sofia, Daniel, Will, Thiago, Michal, Leigh and Joël, that made this thesis possible in the nick of time. It takes a village.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

> "It's the job that's never started as takes longest to finish."
>
> Samwise Gamgee – *The Lord of the Rings*, J.R.R. Tolkien

The neutrino was first proposed by Wolfgang Pauli in 1930 as a massless, parity-conserving, weakly interacting fermion, to maintain energy conservation in the face of the continuous β-decay energy spectrum measured earlier that century. Originally called "neutron" before the neutron's discovery, Pauli stated "I have done a terrible thing, I have postulated a particle which cannot be detected": the neutrino [4]. In the decades since, however, the re-branded neutrino was measured by F. Reines and C. L. Cowan [5] from a nuclear reactor, and later found to be parity-violating, massive, and to exist in (at least) three flavours which oscillate between one-another as they propagate. A whole field of particle physics has developed around these "barely existing" particles, skirting the boundaries of the Standard Model. Since then, such neutrino oscillation has been proven, and most of the involved parameters determined to varying levels of accuracy [6], as the age of high precision neutrino physics is dawning.

However, questions remain to be answered and measurements to be improved, such as the extent of neutrinos' charge-parity (CP) violation or their mass ordering. One such puzzle is the neutrino masses being several orders of magnitude smaller than any other known massive particle. As the only electrically neutral fermion in the Standard Model, Ettore Majorana proposed a novel mechanism for the neutrino to gain mass, which would make them Majorana fermions – unlike all the others being Dirac fermions. If this were true, the small mass could be explained by a so-called "see-saw" mechanism, which would imply physics beyond the Standard Model, and produce lepton number violating processes with far-reaching consequences [4].

The SNO+ detector was thus designed to prove the neutrino's Majorana nature by measuring a process called neutrinoless double-beta decay ($0\nu\beta\beta$), or at the least set limits on this decay's half-life. SNO+ is the inheritor of the Sudbury Neutrino Observatory (SNO), seated 2 km underground in an active mine in the Canadian Shield. It is a large liquid scintillator detector achieving the incredibly low consomogenic and radioactive background levels needed for such a sensitive measurement [7]. It also happens to be ideally situated relative to the nuclear reactors in Ontario to measure their electron antineutrinos undergo neutrino oscillation. SNO+ was already the first to measure reactor antineutrinos in water [8], before being filled with liquid scintillator. A preliminary measurement of neutrino oscillation was then carried out in the so-called "partial-fill" phase [9] (in the process of publication), where the detector was filled with half ultra-pure water and half liquid scintillator. An improved measurement in the "full-fill" scintillator phase is now also approaching publication, where all the water has been replaced with liquid scintillator. It is this measurement with which the present analysis concerns itself. Indeed the result is largely based on work presented herein. As such, while the text is kept as general as possible to serve as a useful resource for all future reactor antineutrino – and related – analyses at SNO+, it focuses, where necessary, on the particular full-fill dataset to be published: 134.5 days of live-time from 30/04/2022 to 05/03/2023.

The presented work is separated into different chapters, flowing with a progressing logic, where possible. First, a brief but fundamental overview of the theory of neutrino physics in the Standard Model is laid out, focusing on how the Majorana neutrino arises, and how it leads to $0\nu\beta\beta$. The emergence of neutrino oscillation from the Standard Model – whether Majorana or Dirac – is also described, though stopping short of any phenomenology which is saved for a later chapter. The SNO+ detector is then presented, with its overall design, novel Tellurium loading technique, its data acquisition and simulation, and calibrations. While these are all primarily motivated by a $0\nu\beta\beta$ measurement, they also enable reactor antineutrino detection, needed for the analysis presented herein. With the context laid, the author's work starts to be presented with the calculation of the expected reactor antineutrino flux, and a detailed uncertainty analysis. A chapter is then dedicated to neutrino oscillation modifying this flux. It focuses on neutrino phenomenology, describing the historical and current best measurements of various oscillation parameters, and the sensitivity of SNO+ to the so-called long-baseline ones. It also covers a novel algorithm by the author, to compute neutrino oscillation in constant matter density, achieving high computation speeds without sacrificing accuracy [10]. Next, the expected backgrounds to this reactor antineutrino signal are reviewed, with their rates, spectra and uncertainties computed for the relevant dataset. The cuts and tagging employed to reduce these to manageable levels are also described. A separate chapter is dedicated to a new classifier, achieving a high degree of purification from the most troublesome remaining background: ^{13}C(α, n)^{16}O events, expected to appear in a technical paper in the near future and to be an important part of future analyses. Finally, the oscillation analysis itself is presented, reviewing all the models in play, a new fitting framework developed for the task, and finally the results of the analysis.

Chapter 2

Neutrino Theory

> "The wide world is all about you: you can fence yourselves in, but you cannot forever fence it out."
>
> Gildor – *The Lord of the Rings*, J.R.R. Tolkien

2.1 Introduction

An overview of the theory of neutrino physics is presented here, using the standard model (SM) of particle physics as a starting point. Beyond the SM features are then explored, showing how they lead to neutrino masses and oscillation, and potentially neutrinoless double-beta decay. Implications of these are only briefly discussed throughout. In this way, the theoretical motivation for the SNO+ experiment is touched upon, while laying the groundwork for neutrino oscillation phenomenology and neutrino-matter interactions, all three being covered in subsequent chapters.

2.2 Neutrinos in the Standard Model

In the SM, after the electroweak symmetry breaking, the three generations of neutrinos (ν_e, ν_μ, ν_τ) are massless left-handed (LH) chiral spinor fields (ψ_L), only interacting via the weak force

$$\mathcal{L} = i \sum_{\alpha=e,\mu,\tau} \left(\bar{\nu}_{\alpha L} \not{\partial} \nu_{\alpha L} + g_Z \bar{\nu}_{\alpha L} \not{Z} \nu_{\alpha L} + \frac{g_W}{\sqrt{2}} \bar{l}_{\alpha L} \not{W}^- \nu_{\alpha L} + \frac{g_W}{\sqrt{2}} \bar{\nu}_{\alpha L} \not{W}^+ l_{\alpha L} \right), \qquad (2.1)$$

where $l_e = e^-$, $l_\mu = \mu^-$, $l_\tau = \tau^-$ are the associated charged leptons. The equations of motion of all three are thus described by three independent massless Dirac equations

$$\not{\partial} \nu_{\alpha L} = 0, \quad \alpha \in \{e, \mu, \tau\}, \qquad (2.2)$$

to zeroth order in the small coupling constants g_Z and g_W. Perturbatively adding in the higher order interactions from the Z and $W^\pm$ bosons leads to neutral-current (NC) and charged-current (CC) interactions respectively. However, these interactions are only ever among either same-flavour neutrinos, or same-flavour neutrinos and charged leptons. The equations of motion for the three flavours are clearly decoupled in this case.

2.3 Beyond the Standard Model

However, the phenomenon of neutrino oscillation was demonstrated by baseline-dependent disappearance of atmospheric muon neutrinos with the Super-Kamiokande experiment [11, 12], and by the flavour transition of solar neutrinos observed by the SNO collaboration. This suggested two things: neutrinos must have mass in order to evolve, and there must be neutrino flavour mixing. More on the history of these measurements is presented in chapter 5. Meanwhile, these two changes can be realised in either one of two straightforward additions to the SM.

2.3.1 Dirac Fermions

The first method is adding a right-handed (RH) neutrino field for each flavour and constructing a Yukawa term, which reduces to a Dirac mass term after the electroweak symmetry breaking. This leads to flavour mixing in exactly the same way as for quarks and the Cabbibo-Kobayashi-Maskawa (CKM) matrix, since diagonalising the mass terms leaves all other terms invariant, except for the weak CC terms

$$\begin{aligned}
\mathcal{L}_D &= \sum_{f=1,2,3} \left(\bar{\nu}_f i \not{\partial} \nu_f + m_D \bar{\nu}_f \nu_f + \frac{g_Z}{2} \bar{\nu}_f i \not{Z} \left(1-\gamma^5\right) \nu_f \right) \\
&\quad + i \frac{g_W}{2\sqrt{2}} \sum_{\alpha=e,\mu,\tau} \left(\bar{l}_\alpha \not{W}^- \left(1-\gamma^5\right) \nu_\alpha + \bar{\nu}_\alpha \not{W}^+ \left(1-\gamma^5\right) l_\alpha \right), \\
\nu_\alpha &= \sum_{f=1,2,3} U^*_{\alpha f} \nu_f,
\end{aligned} \tag{2.3}$$

so that Greek indices represent the flavour eigenstates (defined according to the associated charged leptons), and Latin indices represent the mass eigenstates. m_D is the Dirac mass, U is the (unitary) Pontecorvo-Maki-Nakagawa-Sakata (PMNS) mixing matrix, and $\nu = \nu_L + \nu_R$ and $P_L \psi = \frac{1}{2}\left(1-\gamma^5\right)\psi = \psi_L$ allows one to introduce the so-called V-A structure above. V-A structure will be further explained in chapter 4. Focusing on the PMNS matrix, an $N \times N$ unitary matrix can be parametrised as [4]

$$\begin{aligned}
&U = D\left(\boldsymbol{w} - \boldsymbol{\phi}\right) \left[\prod_{a<b} W^{ab}\left(\theta_{ab}, \eta_{ab} + \phi_a - \phi_b\right) \right] D\left(\boldsymbol{\phi}\right), \\
&W^{ab}\left(\theta_{ab}, \eta_{ab}\right) = \mathbb{1} + \left(cos\theta_{ab} - 1\right)\left(A^{aa} + A^{bb}\right) + sin\theta_{ab}\left(e^{i\eta_{ab}} A^{ab} - e^{-i\eta_{ab}} A^{ba}\right), \\
&[A^{ab}]_{rs} = \delta_{ar}\delta_{bs}, \\
&D\left(\boldsymbol{\phi}\right) = \text{diag}\left(e^{i\phi_1}, e^{i\phi_2}, ..., e^{i\phi_N}\right),
\end{aligned} \tag{2.4}$$

where θ_{ab}, η_{ab}, w_i and ϕ_i are the N^2 free real parameters. ϕ_a and ϕ_b can be chosen so as to cancel out with $N-1$ of the $\frac{1}{2}N\left(N-1\right)$ η_{ab} factors, leaving $1+\frac{1}{2}\left(N-3\right)$ independent. Next, noting that $\bar{\nu} i \not{\partial} \nu = \bar{\nu}_L i \not{\partial} \nu_L + \bar{\nu}_R i \not{\partial} \nu_R$ and $m_D \bar{\nu}\nu = m_D\left(\bar{\nu}_L \nu_R + \bar{\nu}_R \nu_L\right)$, the neutrino and charged lepton fields can be redefined so as to absorb the diagonal phase matrices $D\left(\boldsymbol{\phi}\right)\nu_L \to \nu_L$, $D\left(\boldsymbol{\phi}\right)\nu_R \to \nu_R$ (and likewise for charged leptons and $D^\dagger\left(\boldsymbol{w} - \boldsymbol{\phi}\right)$), without changing the Lagrangian. Therefore, the PMNS matrix, for Dirac neutrinos, can be parametrised according to $W^{ab}\left(\theta_{ab}, \eta_{ab}\right)$

only, which for a 3×3 matrix can be written as

$$U^D = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{13}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{13}} & 0 & c_{13} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{2.5}$$

depending only on three mixing angles $c_{ij} = \cos\theta_{ij}$, $s_{ij} = \sin\theta_{ij}$ and one complex (CP-violating) phase δ_{13}. The redefinition of the fields above effectively eliminated five complex phases (all three w_i's, and two of the three η_{ab}'s), leaving one independent complex phase. Note that the choice of parametrisation is rather arbitrary (particularly which complex phase to keep), but the one above is the most widely used, resembling a rotation matrix. Nevertheless, combining these together, this is

$$U^D = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{13}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{13}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{13}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{13}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{13}} & c_{23}c_{13} \end{pmatrix}, \tag{2.6}$$

clearly leading to a mismatch of flavour and mass states.

This "classic" addition to the standard model is thus the most straightforward. However, it does not explain why the neutrino masses turn out to be so much smaller than any other mass, and it requires adding three extra particles that can otherwise not interact or be measured in any way: right-handed neutrinos.

2.3.2 Majorana Fermions

The second straightforward way of giving the neutrinos mass and flavour mixing was first proposed by Majorana in 1937 [4, 13]. One would think that both LH and RH components are needed to produce a mass term, just like all the other fermions of the SM. However, this turns out to be false for an electrically neutral fermion. To see this, first one defines a spinor ψ^C as the charge-conjugate of ψ

$$\begin{aligned} \psi(x) &\xrightarrow{C} \psi^C(x) = \xi_C \mathcal{C}\overline{\psi}^T(x), \\ \overline{\psi}(x) &\xrightarrow{C} \overline{\psi^C}(x) = -\xi_C^* \psi^T(x)\mathcal{C}^\dagger, \end{aligned} \tag{2.7}$$

with $\mathcal{C}$ the charge conjugation operator, obeying

$$\begin{aligned} &\mathcal{C}^\dagger\mathcal{C} = \mathcal{C}\mathcal{C}^\dagger = \mathbb{1},\ \mathcal{C}^T = -\mathcal{C},\ \mathcal{C}\left(\gamma^\mu\right)^T\mathcal{C}^\dagger = -\gamma^\mu, \\ &\mathcal{C}\left(P_{L/R}\right)^T\mathcal{C}^\dagger = P_{L/R},\ |\xi_C|^2 = 1. \end{aligned} \tag{2.8}$$

ξ_C is the charge conjugation phase associated to each field, and is often ignored. Recall that the spinors can be written in terms of their Fourier decomposition in the framework of Quantum Field Theory (QFT) as

$$\hat{\psi}(x) = \int \frac{d^3\vec{p}}{(2\pi)^3 2E} \sum_{s=1,2} \left(\hat{a}_s(\vec{p}) u_{(s)}(p) e^{-ip\cdot x} + \hat{b}_s^\dagger(\vec{p}) v_{(s)}(p) e^{ip\cdot x} \right), \tag{2.9}$$

where $u_{(s)}(p)$ and $v_{(s)}(p)$ are spinor fields associated with spin s. $\hat{a}_s(\vec{p})$ and $\hat{b}_s^\dagger(\vec{p})$ are the associated particle annihilation and anti-particle creation operators respectively, since

$$|(p_1, s_1), ..., (p_n, s_n); (q_1, r_1), ..., (q_m, r_m)\rangle = \hat{a}_{s_1}^\dagger(\vec{p}_1)...\hat{a}_{s_n}^\dagger(\vec{p}_n)\hat{b}_{r_1}^\dagger(\vec{q}_1)...\hat{b}_{r_m}^\dagger(\vec{q}_m)|0\rangle, \tag{2.10}$$

for an n-particle, m-anti-particle state (remembering these operators all anti-commute with each other), and $|0\rangle$ is the vacuum state. Notice that performing the charge conjugation (2.7) effectively flips $\hat{a} \leftrightarrow \hat{b}$, switching particles with antiparticles and vice-versa. Therefore ψ^C is the antiparticle of ψ. Now, setting the *Majorana condition*

$$\psi^C = \psi, \tag{2.11}$$

one can construct the exact same Lagrangian as for the Dirac neutrinos (2.3), with the constraint $\nu_R = \nu_L^C$. This reduces the four independent components of ν to two independent components, like for a Weyl spinor, and the two coupled Dirac equations to one $i\partial\!\!\!/ \nu_L = m_M \nu_L^C$. The same diagonalisation can thus be performed, leading to flavour mixing for the weak CC sector, given by the PMNS matrix. However, it turns out this matrix cannot be parametrised in exactly the same way. Recall the diagonal phase matrices in (2.4) were eliminated by redefining $D(\boldsymbol{\phi})\,\nu_L \to \nu_L$, $D(\boldsymbol{\phi})\,\nu_R \to \nu_R$, and similarly for the charged leptons. For Majorana neutrinos this is not possible since the first condition implies $D(\boldsymbol{\phi})^\dagger\,\nu_L^C \to \nu_L^C$, which clearly does not leave the Majorana mass term invariant

$$\frac{m_M}{2}\nu = \frac{m_M}{2}\left(\overline{\nu}_L \nu_L^C + \overline{\nu_L^C}\nu_L\right). \tag{2.12}$$

The extra factor of $\frac{1}{2}$ is due to the two fields not being independent anymore, so that the equations of motion are the same. Also note that ν_L^C is RH and transforms like ν_R under Lorentz transformations, so that the Lagrangian is a Lorentz scalar, as required. Back to the parametrisation, the inability to rephase the fields correctly means that only the w_i's can be absorbed (by the charged leptons, as before). Therefore, there are two extra complex phases (λ_1, λ_2) in the Majorana parametrisation of the PMNS matrix compared to the Dirac case, which one can choose as

$$U^M = U^D D^M, \quad D^M = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\lambda_1} & 0 \\ 0 & 0 & e^{i\lambda_2} \end{pmatrix}. \tag{2.13}$$

It will turn out that this difference has no impact on neutrino oscillations, which are covered in the next section. First however, note the mass term above (2.12) must be the post symmetry-breaking form of a term invariant under the SM's $SU(3) \times SU(2) \times U(1)_Y$ symmetries, just like the Yukawa terms are for the Dirac mass term. It turns out that the Majorana constraint again means that the weak hypercharges of every field are already set to satisfy other conditions, so no combination of fields with energy dimension of 4 ($[E]^4$) or less can produce such a term (recalling that spinors $\sim [E]^{3/2}$ and scalars $\sim [E]$). Any terms that do not satisfy this end up being non-renormalisable [4, 14]. Therefore, the presence of a Majorana mass term is a clear indication of physics beyond the SM. Just like Fermi theory was a (non-renormalisable) low-energy limit to the weak force, this could be a low-energy limit to a larger theory. For example, the lowest dimension ($[E]^5$) term one can construct, that satisfies the SM symmetries, is (with only one flavour for simplicity)

$$\mathcal{L}_5 = \frac{g}{\mathcal{M}} L_L^T \sigma_2 \phi \mathcal{C}^\dagger \phi^T \sigma_2 L_L + \text{h.c.}, \tag{2.14}$$

where h.c. means the Hermitian Conjugate of the previous term, σ_2 is the second Pauli matrix, $L_L = (\nu_L, l_L)$ and $\phi = \frac{1}{\sqrt{2}}(\varphi, v + H)^T$ is the scalar field that "contains" the Higgs field H. $\mathcal{M}$ must have dimensions of energy so that $\mathcal{L} \sim [E]^4$, and just as for Fermi theory ($G_F = \sqrt{2}\,(g_W/2M_W)^2$), it may be of the energy scale of the unification into the higher theory, which could explain the smallness of the neutrino mass $m_M = \frac{gv^2}{\mathcal{M}}$. This is a type of see-saw mechanism, one of a plethora of schemes to explain the small mass [15].

2.4 Neutrino Oscillations

Whether the neutrino masses and mixing are generated by the Dirac or Majorana frameworks, they lead to the effect known as neutrino oscillation, in the following way. Neutrinos are produced in CC interactions in pure flavour states $|\nu_\alpha\rangle,\ \alpha \in \{e, \mu, \tau\}$, which as was established above, are composed of a superposition of the mass states $|\nu_k\rangle,\ k \in \{1, 2, 3\}$

$$|\nu_\alpha\rangle = \sum_k U^*_{\alpha k} |\nu_k\rangle. \tag{2.15}$$

This is assuming the mass differences between the mass states are negligible, so that the kinematics does not constrain some states over others. The normalisation conditions are also chosen as $\langle\nu_k|\nu_j\rangle = \delta_{kj}$, which leads to $\langle\nu_\alpha|\nu_\beta\rangle = \delta_{\alpha\beta}$, from the unitarity of U. In a vacuum, it is the mass states that are eigenstates of the free Hamiltonian ($\hat{H}_0$), and so whose evolution can be computed

$$\partial^\mu |\nu_k(x)\rangle = \hat{P}^\mu_0 |\nu_k(x)\rangle = -ip^\mu_k |\nu_k(x)\rangle, \tag{2.16}$$

where $\hat{P}^\mu_0$ is the (free) spacetime translation operator ($\hat{P}^0_0 = \hat{H}_0$). Assuming plane-wave solutions, this is solved with

$$|\nu_k(x)\rangle = e^{-ip_k \cdot x} |\nu_k\rangle, \tag{2.17}$$

where $|\nu_k\rangle \equiv |\nu_k(0)\rangle$. Now consider a neutrino created in a pure flavour state α at an event $x^\mu = \left(0, \vec{0}\right)$:

$$|\nu_\alpha(x)\rangle = \sum_k U^*_{\alpha k} e^{-ip_k \cdot x} |\nu_k\rangle, \tag{2.18}$$

reusing (an inverted version of) (2.15) to write the RH side in terms of flavour states, and multiplying on the left by $\langle\nu_\beta|$, gives the transition amplitude $\langle\nu_\beta|\nu_\alpha(x)\rangle$. Therefore, the transition probability of detecting flavour state β from an initial production of flavour state α is $P_{\nu_\alpha \to \nu_\beta}(x) = |\langle\nu_\beta|\nu_\alpha(x)\rangle|^2$, or explicitly [4, 11]

$$P_{\nu_\alpha \to \nu_\beta}(x) = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} e^{-i(p_k - p_j) \cdot x}. \tag{2.19}$$

Now, due to the extremely small neutrino mass, one can approximate that for neutrinos with detectable energies, $t = L$ (the distance travelled), and so

$$p_k \cdot x \;=\; E_k t - \vec{p}_k \cdot \vec{x} \;=\; (E_k - p_k)\, L \;=\; \frac{E_k^2 - p_k^2}{E_k + p_k} L \;=\; \frac{m_k^2}{E_k + p_k} L \;\approx\; \frac{m_k^2}{2E} L, \tag{2.20}$$

so that as a result, the transition probabilities depend on the mass differences $\Delta m^2_{kj} \equiv m^2_k - m^2_j$

$$P_{\nu_\alpha \to \nu_\beta}(L,E) = \sum_{k,j} U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \exp\left(-i \frac{\Delta m^2_{kj} L}{2E}\right). \tag{2.21}$$

Finally, using the unitarity of U (which means $P_{\nu_\alpha \to \nu_\beta}[k \leftrightarrow j] = P^*_{\nu_\alpha \to \nu_\beta}$), and $P_{\nu_\alpha \to \nu_\beta}(0) = \delta_{\alpha\beta}$, one can arrive at the common and useful form

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta}(L,E) = &\delta_{\alpha\beta} - 4\sum_{k>j} \Re\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) \\ &+ 2\sum_{k>j} \Im\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin\left(\frac{\Delta m^2_{kj} L}{2E}\right). \end{aligned} \tag{2.22}$$

For antineutrinos, the PMNS matrix elements in (2.15) are simply complex conjugated, meaning the same derivation can be performed, and only the imaginary component changes sign

$$\begin{aligned} P_{\overline{\nu}_\alpha \to \overline{\nu}_\beta}(L,E) = &\delta_{\alpha\beta} - 4\sum_{k>j} \Re\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right) \\ &- 2\sum_{k>j} \Im\left[U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}\right] \sin\left(\frac{\Delta m^2_{kj} L}{2E}\right), \end{aligned} \tag{2.23}$$

where for Majorana neutrinos, positive helicity ones act like Dirac antineutrinos, while negative helicity ones like Dirac neutrinos. Therefore, when antineutrinos are mentioned, it is conventionally understood to refer to either Dirac antineutrinos or positive helicity Majorana neutrinos, which are thus both described by the latter equation (2.23) [4]. For survival probabilities (where $\alpha = \beta$), the imaginary term obviously vanishes, leaving

$$P_{\nu_\alpha \to \nu_\alpha}(L,E) = 1 - 4\sum_{k>j} |U_{\alpha k}|^2 |U_{\alpha j}|^2 \sin^2\left(\frac{\Delta m^2_{kj} L}{4E}\right), \tag{2.24}$$

applying equally to neutrinos and antineutrinos.

Notice that all these probabilities are unaffected by the presence of the extra Majorana phases D^M in the PMNS matrix parametrisation, since under rephasing $U_{\alpha k} \to e^{i\phi_\alpha} U_{\alpha k} e^{i\theta_k}$,

$$U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j} \to U^*_{\alpha k} U_{\beta k} U_{\alpha j} U^*_{\beta j}, \tag{2.25}$$

is invariant. Therefore, oscillation experiments cannot be used to distinguish Dirac and Majorana neutrinos – as stated previously – even accounting for matter effects. This will be demonstrated in chapter 5, where oscillation in constant matter density, and specific application to long baseline $\overline{\nu}_e$ survival probability will be covered in detail.

On a final note, a few subtle assumptions were made and details ignored in this derivation, such as the correspondence of the production and detection states with flavour states, the negligible impact of mass difference on the created mass states, and their different group velocities and propagation as wavepackets being negligible. These turn out to either be good approximations, or unimportant in oscillation experiments, but to show this a full derivation using wavepackets in the framework of QFT, rather than plane waves, must be used. See for example [4, 16, 17, 18, 19]. Essentially, there is a localisation term that is unimportant due to the detector region being much smaller than the oscillation wavelength, and a coherence length due to the mass eigenstates decohering from their different group velocities. This can also be ignored for our purposes ($L \ll L_{\text{coh}}$).

2.5 Neutrinoless Double Beta Decay

The Dirac and Majorana parametrisations of the PMNS matrix cannot be distinguished via oscillation experiments, and both types are equivalent kinematically since their masses are not constrained. Therefore the only way to test this is via direct evidence of the lepton number violating Majorana interactions, introduced by the associated mass term (2.12). More simply, since they are their own antiparticles, they can annihilate, leading to processes such as neutrinoless double beta ($0\nu\beta\beta$) decay. Standard $\beta^{\pm}$ decays occur when an atomic nucleus can lower its energy by changing its atomic number by $\Delta Z = \pm 1$, so that a proton is converted to a neutron, or vice-versa, with the corresponding electron (positron) - antineutrino (neutrino) pair being emitted. However for some nuclei, changing Z by only ± 1 raises its energy, due to the difference between even and odd pairing energies, while the rarer change by $\Delta Z = \pm 2$ simultaneously can still sometimes lead to a more stable (lower energy) configuration. This is standard double beta decay ($2\nu\beta\beta$) [20]

$$(Z, A) \rightarrow (Z \pm 2, A) + 2e^{\mp} + 2\overset{(-)}{\nu}_e, \tag{2.26}$$

where the only reliably detectable products are the electrons (positrons), since the recoil on the nucleus is negligible, and neutrinos are obviously hard to detect. Therefore the detectable products exhibit a continuous energy spectrum dependent on the energy taken away by the neutrinos. Now, if neutrinos are Majorana fermions, the two produced by this process can annihilate in the associated Feynman diagram, shown in figure 2.1. This is $0\nu\beta\beta$, and means that effectively all the energy is carried away by the electrons. If this process occurs, one would expect a detection peak right at the highest energy of the energy spectrum - at the Q-value of the process [7].

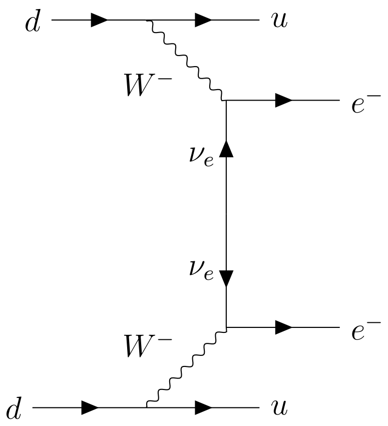

Figure 2.1: Neutrinoless double beta decay.

$$(Z, A) \to (Z + 2, A) + 2e^-, \tag{2.27}$$

For $2\nu\beta\beta$ decays to be measurable, isotopes are needed where this is possible, while single β decay is forbidden or heavily suppressed. These are nuclei with even-even → even-even transitions, and 20 such isotopes have been measured to undergo $2\nu\beta\beta$ [21, 20]. Due to its high natural abundance, ^{130}Te was chosen as the target isotope for SNO+. As such, a low energy detection threshold, high energy resolution, and low backgrounds are required to measure its 2.5 MeV decays, described in the next chapter. These requirements also make the reactor antineutrino measurements, the focus of this analysis, possible.

Chapter 3

The SNO+ Experiment

> "I am old, Gandalf. I don't look it, but I am beginning to feel it in my heart of hearts."
>
> Bilbo Baggins – *The Lord of the Rings*, J.R.R. Tolkien

3.1 Introduction

SNO+ is a large underground neutrino detector, based on the original Sudbury Neutrino Observatory (SNO), which operated between 1999 and 2006 [20, 22, 7]. Its measurments of ^{8}B solar neutrinos helped resolve the solar neutrino problem, culminating in Arthur McDonald of the SNO experiment being jointly awarded the 2015 Nobel Prize in Physics, along with Takaaki Kajita of the Super-Kamiokande experiment [23, 24]. To achieve the primary goal of detecting neutrinoless double-beta decay ($0\nu\beta\beta$), SNO+ has been upgraded and changed from the original detector: primarily a change of detector medium from heavy water to tellurium-doped liquid scintillator. This allows for physical events, such as the ^{130}Te's double beta decays, to be measured down to lower energies and with higher energy resolution, while sacrificing directionality information.

A general overview of the design of SNO+ is presented, followed by a focus on its novel Te-loading method, and its data acquisition system. A brief explanation is then given about the data quality checks put in place, and some of the calibration systems at SNO+. Finally a few of the physics goals enabled by the experiment are listed.

3.2 Overall Design and Upgrades from SNO

SNO+'s primary objective is the detection of neutrinoless double-beta decays ($0\nu\beta\beta$) [7]. To this end, the SNO detector was re-used, being based at SNOLAB in Sudbury, Canada, in an active mine with 2070 m of flat rock overburden to shield against cosmic muons. It is composed of a large 6 m radius acrylic vessel (AV) surrounded by photomultiplier tubes (PMTs), and filled with liquid organic scintillator (specifically Linear Alkylbenzene (LAB) doped with 2 gL^{-1} of PPO (2,5-diphenyloxazole) wavelegnth shifter), which is non-polar, allowing purification of up to 1000 times less radio-impurities than water [20, 23, 25, 26]. Additionally, scintillators produce far more photons per energy deposited, leading to about two orders of magnitude more light in the PMT sensitivity region, which translates to a lower trigger threshold (sub-MeV) and a higher energy resolution [27, 28]. More details on scintillator physics are presented in chapter 7. This is a big change from the predecessor SNO, which relied only on the Cherenkov light produced by charged particles moving through its heavy water. Cherenkov light is also produced in liquid scintillator, but is completely dominated by the scintillation light. ^{130}Te will be added to the scintillator mix in the near future to actually (and hopefully) produce the $0\nu\beta\beta$ decays. This isotope was selected for the experiment due to its long $2\nu\beta\beta$ half-life (7.71×10^{20} years [29], to not overwhelm the $0\nu\beta\beta$ signal), lack of inherent absorption lines in the visible light range (produced by scintillators), and high natural abundance (34%), making enrichment unnecessary [7]. More information on the innovative ^{130}Te purification and deployment technique is provided in the next section. Overall this setup enables the very low backgrounds and large scalability necessary to attempt a $0\nu\beta\beta$ measurement.

In order to deploy this particular Te-loaded scintillator cocktail, the SNO+ detector medium was – and still is being – modified in stages, assessing radiogenic backgrounds and detector performance each time. These lead to different so-called phases, where data taken within one phase has broadly consistent detector conditions. SNO+ was originally filled with ultra-pure water (UPW) during its commissioning, until mid-2019, making up the “water phase”. The water was then gradually replaced with LAB (doped with 0.6 g/L PPO), though the process was halted due to COVID lockdown restrictions, leaving the detector with roughly half UPW and half liquid scintillator. This “partial-fill phase” lasted for most of 2020, until scintillator filling resumed. With filling complete, the “full-fill phase” began in 2021, though the PPO concentration was increased incrementally, reaching the target 2.2 g/L value at the start of 2022. It is this 2.2 g/L PPO full-fill phase that is of interest in this work, and all data is taken from this 2022-2023 period unless otherwise stated. Other additives (bis-MSB and BHT, discussed later) have since been added in mg/L quantities from around mid to end 2023, though they do not enter this work. Tellurium loading itself has not begun as of the writing of this text. Note that the nomenclature of the full-fill phase(s) is currently under discussion and will likely in future work.

The other primary upgrades and developments are hold-down ropes to counteract the buoyancy of the AV due to the scintillator’s lower density with respect to the surrounding water, a scintillator purification system, and an upgrade to the electronics and data acquisition system (DAQ) to deal with the increased event rate.

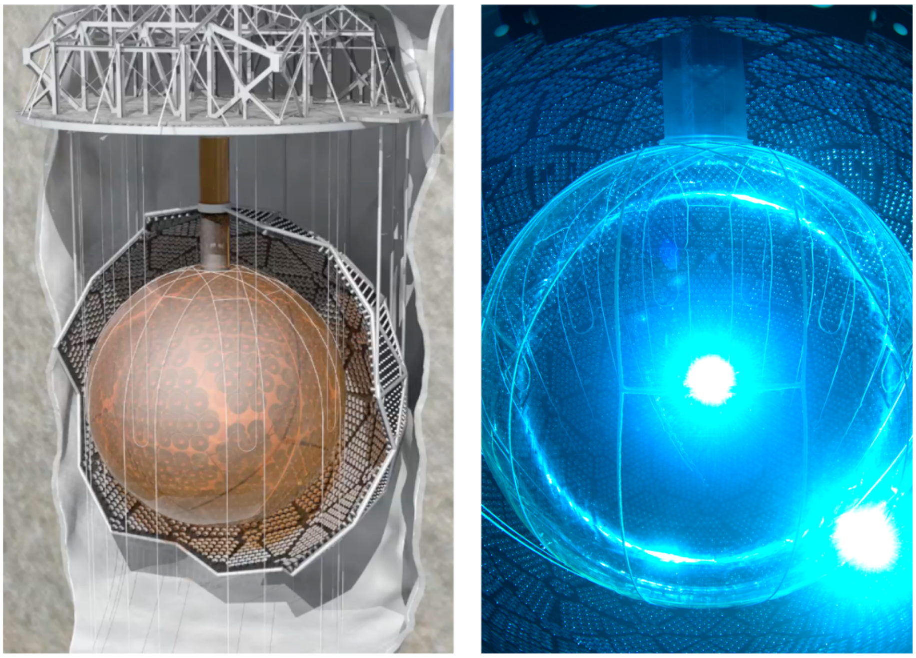

Figure 3.1: The SNO+ detector, as an artistic impression on the left and a photograph from inside the detector on the right [7].

The AV is surrounded by 9362 inward-looking PMTs, each surrounded by a light concentrator to increase coverage, all mounted on a PMT support structure (PSUP) encompassing the AV at about 8.35m from its centre. SNO had about a 54% total coverage, including the concentrator reflectivity, which has since partly reduced for SNO+ due to aging [30]. PMTs can also be taken offline temporarily due to electronics issues. In the dataset of interest, roughly 80-90% of the originally installed PMTs are online. 91 outward looking PMTs (OWLs) are also mounted (without concentrators), to pick up light from incoming muons and other light sources from outside the PSUP. This all sits in 7000 tonnes of ultra-pure water, filling up the cavity outside the AV to shield it from radioactive sources in the PSUP and cavity walls [7, 22, 20]. Other concerns over background levels include radon leaching from the AV walls, since it was exposed to radon-rich air during refurbishment - requiring regular recirculation of the water and scintillator - and further radon contamination from the mine air - mitigated with cover gas systems, flushed periodically with highly purified nitrogen. The backgrounds will be covered in more detail in chapter 6, along with efforts to reduce their impact.

3.3 Tellurium Loading

Another primary advantage of using liquid scintillator is that one can dissolve heavy metals such as ^{130}Te in it, with long term stability and good optical properties [7]. However, the loading process needs to minimize the quantity of isotopes present in the scintillator cocktail with decays around the ^{130}Te Q-value (2.527 MeV). As such, the requirement of $\mathcal{O}(10^{-15})$g/g concentration of uranium (U) and thorium (Th) in the mix was chosen, meaning $< 10^{-13}$ g/g and $< 5 \times 10^{-14}$ g/g of U and Th respectively in the tellurium itself [7, 31].

The tellurium was procured as telluric Acid (TeA) in crystallised form and is stored underground in SNOLAB since 2015 to allow radioactive cosmic ray products to decay away. Even so, the U and Th concentrations still need to be reduced by a factor of roughly 200 and 600 respectively. Two chemical plants were built underground for this purpose: the TeA purification plant and the butanediol (BD) synthesis plant. All wetted surfaces and equipment were also treated to avoid any leaching or contamination. The TeA is purified in the TeA plant by 3 cycles of dissolution in hot UPW, passing through a fine particle filter, recrystallisation (by reducing temperature and pH), and rinsing with clean nitric acid. It is then recrystallised one last time and cleaned with only UPW [7, 32].

The result is then transfered to the BD synthesis plant, where it is heated under vacuum and mixed with 1,2-butanediol, to produce a LAB-soluble product referred to as tellurium butanediol (TeBD). This is mixed with liquid scintillator extracted from the AV at a 1:1 ratio and then transferred to the scintillator plant for further dilution [7].

3.4 Electronics and Data Acquisition

The electronics of the detector are made up of 19 crates, each containing 16 front-end cards (FECs) and 16 PMT interface cards (PMTICs). Four daughterboards (DBs) are attached to each FEC, and contain custom integrated circuits for 8 PMT channels that perform the “important signal processing” [22]. A schematic of the SNO+ electronics is shown in figure 3.2a. Each crate thus deals with the trigger creation and analogue data processing for 512 PMT signals, roughly covering a vertical slice of the detector like the example in figure 3.2b, except for crate 19 which deals only with the 91 OWLs. The PMTs are typically held at a high voltage of around 2000 V, varying from 1800 V to 2350 V across the crates. To deal with the increased data rate from using liquid scintillator, the readout system from the crates was upgraded from that used by SNO, as described in detail in Ref. [7].

The data acquisition system (DAQ) needs to be able to handle the large data-flow caused by photons hitting the PMTs nearly continuously, while only outputting potentially interesting physics in a tractable quantity for further treatment. This is achieved by recording the time and charge of PMT hits, and grouping them into events. Any PMT hit above a certain charge threshold is buffered into its FEC’s memory, while two square waves (one 100 ns and one 20 ns long) are sent to a series of analog master trigger cards (MTC/A+, upgraded from SNO’s MTC/A). These MTC/A+’s sum over the 100 ns and 20 ns pulses from all PMTs detector-wide, and issue a global trigger (GT) when a (programmable) threshold is passed. The 100 ns pulses are the standard signal used for this triggering, called the N100 trigger. The PMT hit time and charge information are then collected from all the FECs, and sent to be digitized as an event with a unique label: its global trigger ID (GTID) [22, 20, 7]. In this way, most low-level random noise is filtered out “at the gate”. Meanwhile, most physical events of interest will produce light that is picked up by PMTs around the whole detector, with very roughly one photon expected to be detected per PMT per event, as will be discussed in section 3.6.2.

The time window of a recorded event is 400 ns long, and is followed by a 420 ns dead-time [2], which is compensated for by an automatic re-trigger of events extending beyond the standard 400 ns [7]. As such, one long event is split into two

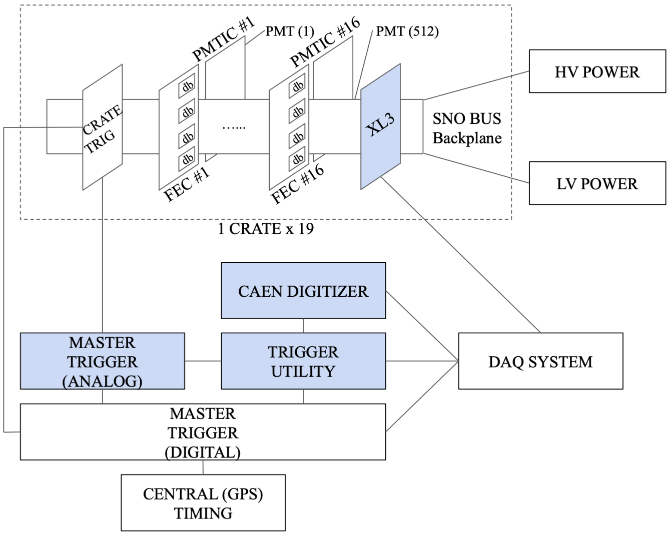

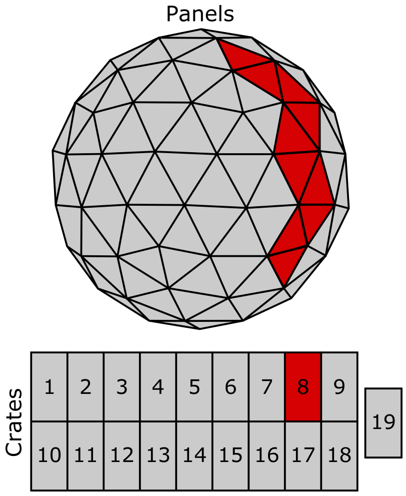

(a) Electronics sketch, where upgrades from SNO are highlighted in grey-blue [7].

(b) An example crate's coverage of the SNO+ detector, highlighted in red. This crate is connected to the PMTs installed on the highlighted panels of the PMT support structure (PSUP), in a vertical slice of the detector. Not accurate, for illustration purposes only.

Figure 3.2: Schematics of the SNO+ detector's electronics layout.

shorter events separated by roughly 400 ns or less.

3.5 Data Quality

Thanks to all the hardware and software deployed at SNO+, it is an incredibly sensitive detector. This, however, means that it is sensitive to signals that are clearly not interesting physics, but effectively noise. It is also very fine-tuned and can be prone to electronics or data-flow mishaps at times. As such, a slew of data quality checks were developed, broadly divided into two categories: data cleaning and run-selection.

Data cleaning consists of a series of checks to flag any events that are clearly not physics events, which are run during the processing of recorded data. In particular, MeV-scale events caused by the detector's so-called "intrumentals" can occur, such as "flashers": the static discharge PMTs, which are held at high voltage, producing a flash of light. Unwanted physics events can also be flagged, such as muons, though their effectiveness is insufficient at the time of writing. The problems posed my muons and how they are mitigated are covered later, in section 6.3.1. Nevertheless, all these "bad events" can then be easily ignored by analysers. This is always performed for any data studied in this analysis.

Run-selection takes a broader view, judging the suitability of an entire run of data; all data is recorded in manageable chunks called runs, which are generally about an hour long, though sometimes less if a run was cut short. Run-level information is collected from a variety of sources. Processors run through all the data to estimate quantities such as the number of bit-flips, whether the internal clocks ran as expected, etc. The state of the detector is an important factor: a minimum number of online crates and PMT coverage are required for instance. Any human activity on deck (just above the detector) deemed disruptive also discounts a run, since the SNO+ detector can be sensitive to extra light, static discharge or physical shocks. Various run-selection criteria are applied to the runs, with different levels of restrictiveness tailored to different analyses. Automation of this process was carried out by the author, with information now displayed online for the SNO+ collaboration. The data used throughout this analysis is a subset of only the highest quality runs: the gold list, which requires all crates to be online for instance. However, work is currently underway to show that less high quality runs can also be included, and checks improved to accommodate this, which is expected to increase the effective livetime in the near future.

3.6 Event Simulation and Reconstruction with `RAT`

The digitised events that pass the above data quality checks are then further processed, in order to reconstruct higher level quantities of the original *physics event* which caused the *measured event*. However, event simulation is first discussed, since both real and simulated events are subject to the exact same reconstruction to ensure consistency between them. Data processing, simulation and analysis is all performed by RAT (Reactor Analysis Tool), which is a software package originally developed for the Braidwood reactor experiment [7], based on the physics simulator GEANT4, and using ROOT classes in `C++` [7, 33, 34].

3.6.1 Event Simulation

RAT simulates the SNO+ setup with data saved in run-specific tables from the custom “RATDB” database, which is updated based on calibrations and other information. Monte Carlo simulations are performed by creating an initial event vertex containing the particle type(s), energy(ies), position(s), direction(s) and time(s), based on some initial model such as the reactor antineutrino flux and inverse beta decay (IBD) interaction. RAT then hands this over to GEANT4, to simulate the propagation and production of photons and secondary particles. For example, IBD events (covered in chapter 4) are generated by RAT, and the resulting positron and neutron are initialised in GEANT4, with appropriate properties. It simulates physics at randomly sampled intervals based on the mean free path of each particle. GEANT4 has base classes to deal with electromagnetic and hadronic physics, as well as for handling particle decays. RAT then takes the emitted photons, and models their interactions with the different detector components (scintillator, AV, PMT, water, etc), whose properties are all calibrated. PMTs are simulated as "grey discs", with position-dependent chances to absorb or reflect photons, depending on their wavelength, and based on wavelength sensitive calibrations of the PMTs’ quantum efficiencies (both in-situ and ex-situ). If a photon is absorbed, the resulting photoelectrons (PEs) are simulated in the pseudo-DAQ, which emulates the process real data goes through, including electronic effects such as noise, finite rise times and trigger windows. Trigger signals are then simulated, producing the same information output as for real data, but with the particle’s Monte-Carlo (MC) history and “true MC” quantities also attached [22, 20, 7].

3.6.2 Event Reconstruction

Reconstruction then works the same for both real and simulated data. The only information available is the total number of recorded PMT hits during the event (N_{hit}), and the relative timing and position of these. All higher level quantities, such as original event’s deposited energy, position and time, must be inferred from this.

Energy

The deposited energy is perhaps the most straightforward quantity to reconstruct. As will be explained in section 7.2.3, the liquid scintillator emits a number of photons roughly proportional to this energy. The number of these photons leads to an expectation of less than one photon hitting each PMT during an event, for the MeV-scale energies of interest. As such, the N_{hit} is roughly proportional to the event’s original energy. More specifically, a ratio of roughly 300 N_{hit} per MeV is observed for electrons with kinetic energy up to around 2 MeV or so. Note that other particles or processes may have different proportionality constants, due to the different amounts of light production induced in the liquid scintillator – an example of an effect known as scintillator quenching, and described further in section 7.2.3. However, energy reconstruction always assumes that of the electron. As such, the reconstructed energy of certain particle types may differ from their true energies, particularly for heavier particles such as α’s. Nevertheless, the vast majority of events detected by SNO+ are well modeled by this assumption, and so these reconstructed quantities are used anyway.

However, this proportionality begins to break down at higher energies, as the number of photons increases, and thus so too does the likelihood of some PMTs being hit by more than one photon. Events occurring closer to the detector’s edge also increase this likelihood, since more photons are concentrated onto fewer PMTs. If more than one photon strikes a PMT in such a short time, the photo-electrons can pile up, recording only one hit where multiple should have been recorded. Therefore, the N_{hit} of higher energy events can be lower than one would expect. A way around this is to observe the statistical properties of ensembles of similar PMTs. The expected number of times a single PMT was hit can be determined from the ratio of the mean number of recorded PMT hits per PMT to the total number of recorded PMT hits in that ensemble. Finding this for every PMT ensemble in the detector allows one to adjust the raw N_{hit} to the corrected N_{hit}, which is the number of PMT hits that “should have been” measured. This is then used in the proportionality relationship described above, enabling energy reconstruction to a good approximation [20].

Position and Time

While the number of PMT hits is used to determine the event energy, the position and time of and event are instead reconstructed from the relative timing of the PMT hits, and where these PMTs are located – the PMT hit pattern.

First, consider a short burst of photons ($\mathcal{O}(1)$ ns) emitted from the scintillator during an event located in the centre of the detector. The resulting distribution of the hit times of each PMT t_{hit} (recall each PMT will almost always record either zero or one hit during an event) thus also exhibits a concentrated burst of PMT hits around the same time, since the time-of-flight t_{TOF} from the event to the PMT is roughly the same for all the photons. This burst, or "peak", is followed by a tail of later PMT hits from various effects such as delayed scintillation light, absorption and reemission, and reflections. The peak of the t_{hit} distribution thus provides information on the event time, based on the time-of flight from the detector's centre to the PMTs. Now, if the exact same event were to happen away from the detector's centre, closer to the detector edge, the peak of PMT hits would be more spread out in time, since some photons have further to travel than others before hitting a PMT. The spread of the t_{hit} peak therefore contains information on the radial distance of the event from the detector's centre, while the PMTs with the earliest hit times indicate which edge the event is closest to. Therefore, one can use the distribution of the PMT hit times to determine the event's time and position.

The specific way this information is gleaned, is by first defining the residual hit time (or time residual) for each PMT hit, via

$$t_{\text{res}}\left(t_{\text{evt}}, \vec{r}_{\text{evt}}\right) = t_{\text{hit}} - t_{\text{evt}} - t_{\text{TOF}}\left(\vec{r}_{\text{evt}}\right), \tag{3.1}$$

where only the PMT hit time t_{hit} is known a priori, while the reconstructed event time t_{evt} and position $\vec{r}_{\text{evt}}$ are unknown quantities that are to be determined. $t_{\text{TOF}}\left(\vec{r}_{\text{evt}}\right)$ is the straight-line time-of-flight of a photon, computed between the event and PMT positions[1]. Now, if the reconstructed event position is correctly chosen – the same as the true event position – the short burst of light described in the previous chapter will give rise to a time residual distribution with a narrow peak, followed a tail at higher t_{res}. This is true no matter where the event took place (so long as the position is reconstructed accurately), since the computed $t_{\text{TOF}}\left(\vec{r}_{\text{evt}}\right)$ will eliminate the spread caused by the differing raw hit times (t_{hit}). In other words, the t_{res} distribution with a correctly reconstruction position will resemble the t_{hit} distribution of a similar but centrally located event (i.e. with $\vec{r}_{\text{evt}} = \vec{0}$), though with some time offset. If, in addition, the event time t_{evt} is also correctly reconstructed, then the t_{res} distribution will begin at around $t_{\text{res}} = 0$ (the peak will be at zero). Bringing these two facts together, the better the reconstruction, the more the t_{res} distribution of any event will tend towards the t_{hit} distribution of a centrally located event with its peak near $t_{\text{res}} = 0$. In other words, it will tend towards the following

$$t_{\text{res, central}}\left(0, \vec{0}\right) = t_{\text{hit, central}} - t_{\text{TOF}}\left(\vec{0}\right), \tag{3.2}$$

where $t_{\text{hit, central}}$ are the raw hit times of a centrally located event, and $t_{\text{TOF}}\left(\vec{0}\right)$ is almost identical for every PMT, since the event is centrally located.

Therefore, a $t_{\text{res, central}}$ PDF was constructed from many event simulated in the detector's centre. Then, for each event, the log-likelihood ratio of $t_{\text{res}}\left(t_{\text{evt}}, \vec{r}_{\text{evt}}\right)$ and $t_{\text{res, central}}$, is maximised by varying t_{evt} and $\vec{r}_{\text{evt}}$ simultaneously. Likelihood ratios

---

[1] This time-of-flight calculation accounts for the the refractive indices of various traversed materials for their effective light speeds, but otherwise assumes a perfectly straight line path from the event to the PMT; no refraction is included.

will be discussed in more detail in chapters 7 and 8. The values that maximise this are the best fit values, and are thus used as the event's reconstructed time and position [2].

Lastly, various goodness-of-fit parameters are also computed [22], to be used in various higher level analyses.

3.7 Calibration

To achieve realistic PMT hit times for a given event, and robust event simulation and reconstruction, both the detector and the RAT code must be appropriately calibrated. To this end, a suite of optical and radioactive calibrations are employed. These rely on naturally occurring internal radioactivity, or either deployed or permanently installed sources, whether optical or radioactive. A brief overview of the systems relevant to the present analysis is given here, particularly as they provide estimates for systematic uncertainties used later on.

3.7.1 Internal Calibration: ELLIE

In order to measure internal properties of the SNO+ detector in a semi-continuous manner, without introducing any foreign objects into it, the ELLIE system was put in place: the Embedded LED/Laser Light Injection Entity. It consists of fibre-optic cables mounted on the PSUP, between the PMTs, which lead back up to racks on deck, connecting to LEDs or lasers. In this way, light can be injected into the detector with known properties, while avoiding any internal contamination.

ELLIE consists of three modules: AMELLIE, SMELLIE and TELLIE – the attenuation, scattering and timing modules, respectively, displayed in figure 3.3. The first two are geared towards measuring the scintillator's optical properties, while the latter aims to perform the timing calibration of the detector's PMTs and associated electronics [7]. Of primary interest here is the TELLIE calibration, whose goal is to ensure that the PMT hit times recorded by the MTC/A+ are corrected to the physical PMT hit times. All event reconstruction relies on this, as does the classifier presented in chapter 7. A TELLIE calibration consists in successively emitting light from all of its 96 calibrated fibres, spread all around the detector, at known times, and recording the hit times of the PMTs in each successive beam spot. The PMT hit pattern of a typical TELLIE run, in which only one fibre is fired, is shown in figure 3.4. By comparing the known light emission times, the measured hit times of the PMTs from different overlapping beam spots, and the calculated light travel times, their timings can be fitted and thus calibrated. The timing resolution achieved at SNO+ with this and other calibrations is on the order of 1 ns, enabling more detailed analysis of events through such means as pulse shape discrimination.

3.7.2 In-Situ Calibration: ^{214}BiPo Events

As will be covered in section 6.2.1, a very small amount of naturally occurring radioactive isotopes are present in the detector, leading to various decay chain products. ^{214}Bi is produced as part of the ^{238}U decay chain, and β decays into ^{214}Po with a Q-value of 3.27 MeV, which then decays to ^{210}Pb via a 7.8 MeV α. While the ^{214}Bi and the ^{210}Pb have half-lives of 19.9 minutes and 22.2 years respectively, the

^{214}Po has a lifetime of only 237 μs. Therefore, the rapid consecutive ^{214}Bi and ^{214}Po decays make for an excellent coincidence pair, in both time and space, referred to as ^{214}BiPo events. The prompt (first event) β has an average energy of 2.2 MeV, while the delayed (second event) α's energy is quenched down to around an apparent 0.8 MeV. Selecting for event pairs that follow this pattern produces a very pure sample of ^{214}BiPo events. The prompt event can thus be used to compare MC simulations of ^{214}BiPos to data, in order to correctly calibrate the energy reconstruction.

However, energy reconstruction is tuned to events within the central 3.5 m fiducial volume (FV), where the main search for $0\nu\beta\beta$ will take place, but becomes poorer further out. For example, the ^{214}BiPo mean prompt energy deviates from 2.2 MeV for events further away from the detector's centre, as shown in figure 3.5a (left). A tool to correct for this in the full-fill phase was thus created by Anthony Zummo [3], by looking for prompt events in the 1.2 to 3.5 MeV range, and delayed events in the 0.6 to 1.1 MeV range, separated by between 2[2] and 1000 μs in time and less than 0.8 m in space. A pure sample of around 60000 event pairs was selected, and used to construct a position-dependent energy correction map. This correction improves the reconstructed energy out to around 5.7 m, for both simulations and data, via the scaling factor maps displayed in figure 3.5b. The data-simulation agreement is thus equivalent for all radial positions up to 5.7 m, for events around 2.2 MeV.

Systematic uncertainties in the energy reconstruction were also established thanks to this study. The residual differences between the prompt energy spectra of data and simulations in the 5.7 m FV were studied, after the above correction was applied. By computing the ratio of the mean energies of data and simulation, the root-mean-squared (rms) of the difference between this and unity at different radii was found to be 1.8%. This can therefore take on the role of an energy scaling uncertainty. Similarly, the rms of the difference between the standard deviations of their energies at different radii was 6.5%. Assuming the energy resolution roughly scales as $\sqrt{E}$ due to the Poissonian statistics of the N_{hit}, and 6.5% is the uncertainty at 2.2 MeV, this provides a systematic uncertainty in the energy resolution of $4.4\% \times \sqrt{E}$.

The prompt β particles were also used to tune the scintillator timing, discussed later in section 7.2.2, with reference to pulse shape discrimination.

3.7.3 External Calibration: the AmBe Source

The SNO+ collaboration uses an Americium-Berylium (AmBe) source, containing a mixture of powdered ^{241}Am and ^{9}Be inside a sealed container, as an external calibration device. When in use, it is deployed in the external water between the PSUP and the AV. In 2005, the neutron rate from the AmBe source was measured by the SNO collaboration to be (67.39 ± 0.73) Hz, which should only have changed negligibly since then. The ^{241}Am has a half-life of 432 years, decaying with a roughly 5.5 MeV α particle, which is then captured by the ^{9}Be $\mathcal{O}\,(0.01\%)$ of the time. This ^{9}B$(\alpha, \text{n})^{12}$C interaction produces a neutron and a ^{12}C, which is created in an excited state 60% of the time, emitting a 4.4 MeV γ to de-excite. Meanwhile, the neutron thermalises (loses almost all its energy by scattering off particles in the medium) and is later captured, usually on a proton, which releases a 2.2 MeV γ as a delayed event [7]. Since the source is deployed externally, only a small fraction of these propagate inside the AV.

---

[2] This lower limit is to cut out the similar ^{212}BiPo events, with a ^{212}Po half-life of 0.3 μs.

However, the rate is high enough that the 4.4 MeV γ and 2.2 MeV neutron capture event pair can be selected by coincidence tagging, just like the ^{214}BiPo events, and used as another data point in energy reconstruction calibration. In particular, most tagged events occur near the AV, where energy needs the most correcting. Such a calibration has not yet been performed in the full-fill phase, though a basic check of the impact of the ^{214}BiPo energy correction at this energy was performed instead. Figure 3.6 shows much better agreement between data and MC at both the 2.2 and 4.4 MeV peaks.

In the 40% of decays where the ^{12}C is produced in the ground state, the neutron carries away most of the energy and can produce its own lower energy prompt event by recoiling with protons. This latter possibility could in principle be used to produce a sample of pure (α, n)-like events, as described in section 6.3.3, though the neutron would need to carry enough energy inside AV before recoiling to be of use. This possibility is currently still under investigation, and future internal deployments of the AmBe source are being planned.

3.8 Other Physics Goals

Since SNO+ is a large neutrino detector, $0\nu\beta\beta$ decay is not its only objective. Particularly since there is a long period of potential physics time before the ^{130}Te is loaded, where the detector was filled with ultra-pure water, and now scintillator. Even after it is loaded, many of these measurements can carry on being performed [7, 20].

For instance SNO+ recently joined the Supernova Early Warning System (SNEWS), as it is capable of detecting a neutrino flux from a supernova in our galaxy. Solar neutrinos have also been measured by SNO+ [36], hearkening back to the original SNO detector. The energy threshold is also low enough to detect geo-neutrinos – antineutrinos produced by radioactive decays in the Earth's interior – providing an additional way to test current geological models [37].

Finally, the SNO+ detector is located roughly 240 km and 340/350 km from its three closest nuclear reactors, which together produce the majority of its antineutrino flux. This provides high sensitivity to long baseline neutrino oscillations in a different way to the KamLAND experiment, as described in chapter 5. Therefore, measuring the $\overline{\nu}_e$ survival probability is of interest, to further constrain the measurements of long baseline oscillation parameters.

3.9 Summary

The SNO+ detector benefits from an array of calibrated hardware and software, producing both data and MC simulations in the form of events composed of their constituent PMT hits. Reconstruction turns this basic information into higher level quantities such as event times, positions and energies, while various data quality filters ensure that only suitable data is studied. In such a way, a plethora of physics analyses are possible, such as the primary goal of $0\nu\beta\beta$ decay, and the focus of the present analysis: reactor antineutrino oscillation.

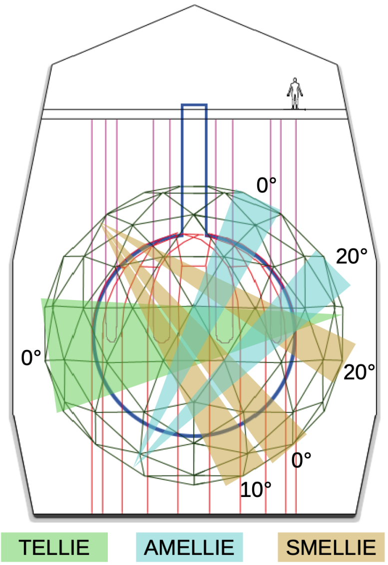

Figure 3.3: Sketch of the three ELLIE systems (TELLIE, AMELLIE, SMELLIE) installed on the SNO+ PSUP (dark green grid), each firing representative beams (shaded triangles) from three of the many injection points into the AV (blue line), either directed at the AV's centre (0°), or slightly off-centre by a particular angle (10° or 20°) [35].

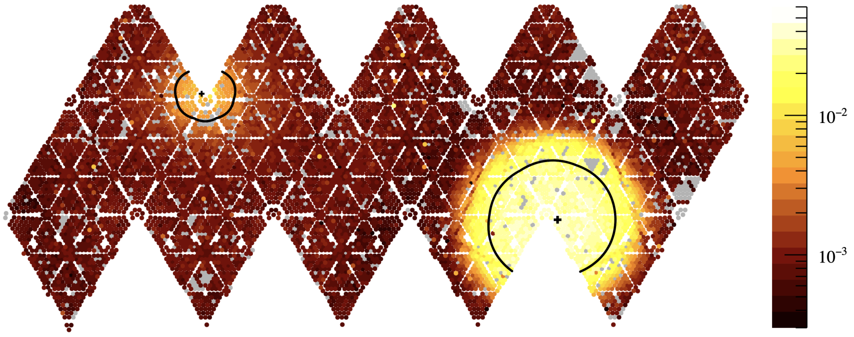

Figure 3.4: An example distribution of PMT hits from a TELLIE run during SNO+'s water phase, showed on a panel flat map, where the panels from figure 3.2b are unfolded to a flat surface. Each dot represents a PMT, with grey ones showing offline PMTs, and the red to yellow scale conveying the relative number of hits each one received during the run. 2D Gaussian distributions are fitted to the beam spot (lower-right) and the near-AV reflection spot (upper left) [7].

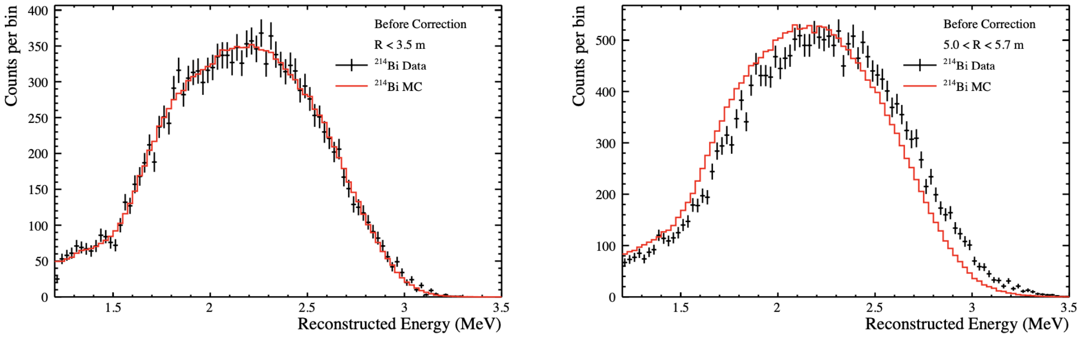

(a) ^{214}BiPo prompt energy spectrum at different radial positions, before position-dependent energy correction. Taken from Ref. [3].

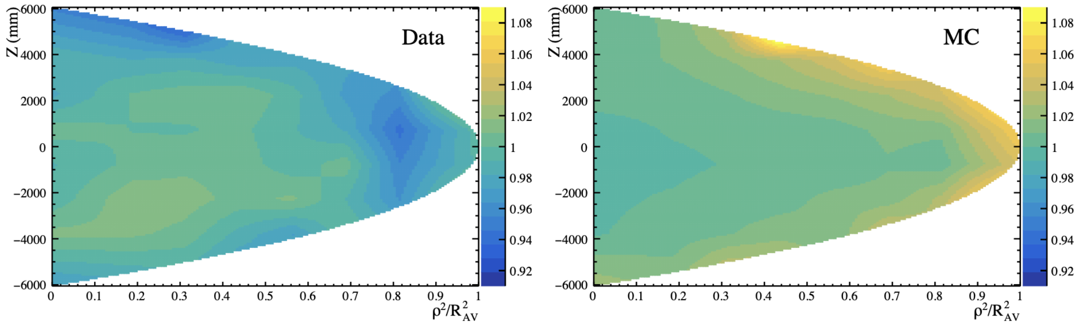

(b) Position dependent energy correction factors, from Ref. [3].

Figure 3.5: Energy reconstruction's position dependence from ^{214}BiPos.

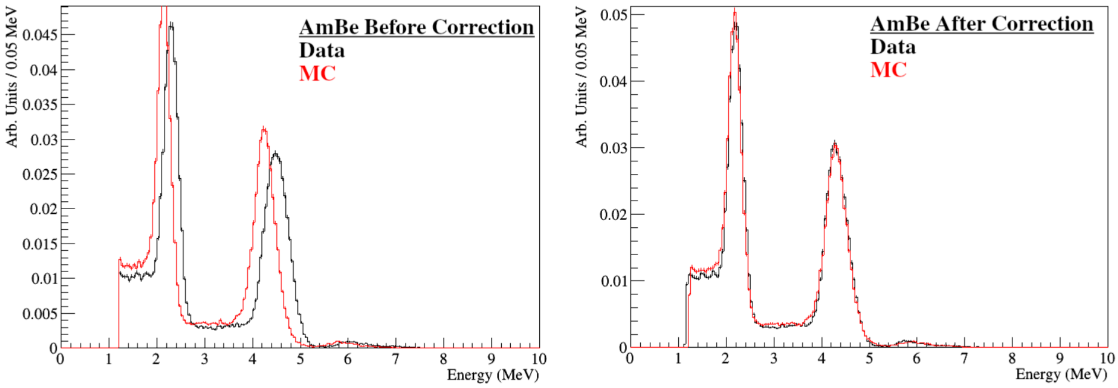

Figure 3.6: Position dependent energy correction's impact on AmBe prompt and delayed events, from T. Kaptanoglu [3].

Chapter 4

Reactor IBD Predictions

> "The dark fire will not avail you, flame of Udûn."
>
> Gandalf – *The Lord of the Rings*, J.R.R. Tolkien

4.1 Introduction

Neutrinos can only interact with other matter via the Weak Nuclear force, as discussed in chapter 2, making their detection a challenge. Inverse beta decays (IBDs) provide a solution to this, being the preferred way of detecting reactor antineutrinos, and in fact enabled the very first detection of any neutrino by F. Reines and C. L. Cowan in 1959 [5, 38, 4]. SNO+ is in an ideal position to measure such reactor antineutrino induced IBDs, for the reasons laid out in this and the following chapters.

An IBD involves the $\overline{\nu}_e$ interacting with a proton and inducing a β^+ decay, as shown in figure 4.1. The signature of this interaction is relatively clear, with an initial prompt event generated by the emitted positron, which induces scintillation light as it slows down and annihilates with an electron. The positron energy is highly correlated with the $\overline{\nu}_e$ energy, so the prompt energy spectrum broadly preserves the features of the $\overline{\nu}_e$ energy spectrum, as will be discussed later. Meanwhile, the neutron thermalises, before being captured by a proton roughly 220μs later, which produces a delayed event at a reliable 2.2 MeV [22].

Now, reactor $\overline{\nu}_e$ are produced with energies of up to around 10 MeV, as will be discussed later. At these energies, only elastic scattering or inverse beta decays (IBD) are likely to occur. Elastic scattering is common to both neutrinos and antineutrinos, while IBDs can only be induced by $\overline{\nu}_e$. Combining this with the fact that nuclear reactors are the largest source of $\overline{\nu}_e$ on Earth at these energies means that detected IBDs are likely produced by reactors. Adding in the clear correlated prompt-delayed signal produced by IBDs allows one to greatly reduce backgrounds, enabling reliable detection of $\overline{\nu}_e$.

Therefore, reactor antineutrinos are used to measure the $\overline{\nu}_e$ survival probability and thus its oscillation parameters. Details on this neutrino oscillation are provided in chapter 5, while its extraction from other background events is covered in chapter 6. The present chapter lays out the calculation of the expected flux and spectrum of reactor produced $\overline{\nu}_e$ IBDs in the SNO+ detector, largely ignoring both oscillation and detection efficiency for now.

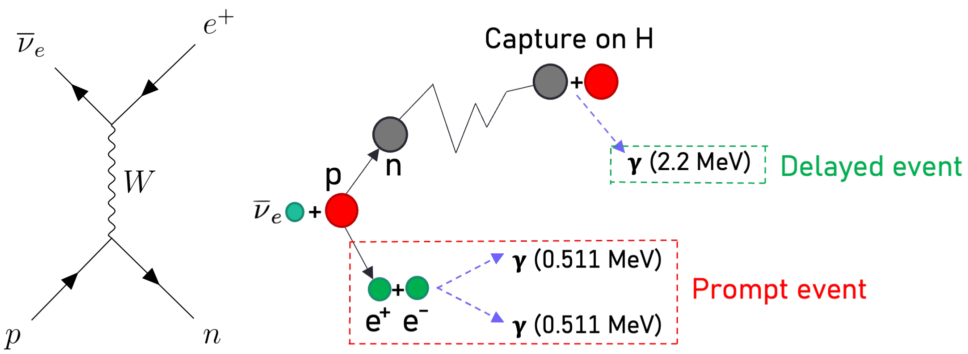

(a) Underlying Process

(b) Event Presentation

Figure 4.1: Inverse Beta Decay: $(Z, A) + \overline{\nu}_e \to (Z-1, A) + e^+$

4.2 Calculation Overview

The way reactor IBD events at SNO+ are predicted is a multi-step process. An overview of the calculation is given here, before a more detailed discussion of each component in subsequent sections. While SNO+ is closest to the three Ontario reactors (Bruce, Darlington and Pickering), at least all the North American reactors must be accounted for in some capacity, to get an accurate accounting of the $\overline{\nu}_e$ flux at SNO+. In this analysis all reactor cores all over the world are included in the calculation for completeness, and will be described in more detail in section 4.3.3. As such, the incoming reactor IBD spectrum is computed by summing over the contributions of all these reactor cores

$$N_{\mathrm{IBD}}(E,t) = \eta(E,t) N_H \sum_{n=1}^{N_{\mathrm{cores}}} \Phi_{(n)}(E,t), \tag{4.1}$$

where N_H is the number of target protons in the detector volume, and $\eta(E,t)$ is the detection efficiency, which may depend on energy and change over time. $\Phi_{(n)}(E,t)$ is the number of IBDs detected at SNO+ from core n, per unit energy, per unit time, per target proton. This is in terms of the "measured" reconstructed prompt energy E, which can be related back to the antineutrino energy E_ν via a convolution

$$\Phi(E,t) = \int_{E_{\mathrm{min}}}^{\infty} \Phi_\nu(E_\nu, t) p\left(E|E_\nu\right) dE_\nu. \tag{4.2}$$

$p\left(E|E_\nu\right)$ is a 2-D PDF which gives the probability density of an antineutrino with energy E_ν leading to an IBD prompt event with reconstructed energy E, normalised to 1 for each energy E_ν. In practice this convolution is realised by generating MC simulations of IBD events from a known antineutrino energy spectrum, and then using the resulting prompt energy spectrum.

One can calculate the input spectrum $\Phi_\nu(E_\nu, t)$ by modelling isotropic emission of antineutrinos from the reactor core, so that

$$\Phi_\nu(E_\nu, t) = P_{\overline{\nu}_e \to \overline{\nu}_e}(E_\nu, L) \frac{\sigma(E_\nu)}{4\pi L^2} F\left(E_\nu, t\right), \tag{4.3}$$

with L the core-detector distance (baseline), and $\sigma(E_\nu)$ the IBD cross-section. $F(E_\nu)$ is the reactor core's emission spectrum: the number antineutrinos emitted per unit energy per unit time. The survival probability $P_{\overline{\nu}_e \rightarrow \overline{\nu}_e}(E_\nu, L)$ is also included here for completeness.

4.3 Outgoing Reactor Core Flux

4.3.1 Introduction

This section concerns itself with the calculation of $F(E_\nu, t)$, which is the number of antineutrinos emitted per unit energy per unit time in a particular reactor core. It is computed in a similar way as for the Daya-Bay experiment [39], via

$$F(E_\nu, t) = \frac{W(t)}{\epsilon} \sum_i f_i n_i(E_\nu)\, c_i^{\mathrm{ne}}(E_\nu, t) + S_{\mathrm{snf}}(E_\nu, t), \tag{4.4}$$

where W(t) is the core's thermal power output over time, while f_i and $n_i(E)$ are the fission fraction and emitted $\overline{\nu}_e$ spectrum for isotope i. The $n_i(E)$ spectra are normalised to one fission, and so integrate to the total number antineutrinos emitted per fission (see figure 4.2 below), while f_i are the fraction of antineutrinos emitted from each isotope out of the total emitted. ϵ is the average energy emitted per fission, computed as a weighted sum of energy emitted per fission for each isotope ϵ_i [40, 41]

$$\epsilon = \sum_i f_i \epsilon_i. \tag{4.5}$$

Finally, $c_i^{\mathrm{ne}}(E_\nu, t)$ models the non-equilibrium effects of long-lived fission fragments, and $S_{\mathrm{snf}}(E_\nu, t)$ is the contribution from spent nuclear fuel rods stored in reactor complexes. The Daya-Bay experiment estimated their contribution from spent fuel to be 0.3%, with an uncertainty of 100%. This is treated as negligible in the present analysis, and thus ignored in what follows. Its contribution should ideally be verified for the nearest reactors in future work.

Daya-Bay also state an average increase of 0.6% to their IBD flux from non-equilibrium effects, with uncertainty of 30% [40], based on calculations in Ref. [42]. This last paper presents a good introduction to the calculation of reactor emission spectra, including the two main methods of obtaining isotope emission spectra $n_i(E_\nu)$: ab initio, and the Institut Laue-Langevin (ILL) method. This second method involves using reference electron spectra, measured over periods of hours to days, to obtain the antineutrino spectra. The non-equilibrium effects arise from discrepancies between spectra obtained this way, and those simulated from more realistic PWR conditions, over burn up cycles on the order of a year [42]. Daya-Bay uses the ILL method, and therefore incorporates the non-equilibrium contributions [40]. However, the analysis presented here will use unfolded spectra from measured reactor antineutrino IBD interactions for its two main contributing isotopes. Combining this fact with the, in any case, small overall contribution of non-equilibrium effects, render them negligible. They are thus also ignored here throughout.

Taking away contributions deemed negligible, the final formula used in this analysis is

$$F(E_\nu, t) = \frac{W(t)}{\epsilon} \sum_i f_i n_i(E_\nu), \tag{4.6}$$

whose components are further discussed below.

4.3.2 Fission Fractions and Isotope Spectra

Reactor Types

Reactor antineutrinos are almost entirely produced by four fissile isotopes – byproducts of the fission process. The four isotopes, with the approximate fraction of emitted antineutrinos from each are ^{235}U (56%), ^{238}U (8%), ^{239}Pu (30%) and ^{241}U (6%) [4]. These fractions are called fission fractions, and their exact values depend both on the type of nuclear reactor, and the evolution of these over time. Addressing the reactor types first, the International Atomic Energy Agency (IAEA) lists six commercial types [43]:

- BWR: boiling water reactor.
- FBR: fast breeder reactor.
- GCR: gas cooled reactor (graphite moderated).
- LWGR: light water (cooled) graphite (moderated) reactor.
- PHWR: pressurised heavy water reactor.
- PWR: pressurised water reactor.

The most common in the world are PWRs, which use enriched uranium as fuel, and a pressurised water circuit to transport heat from the fuel to a separate second water circuit, which then produces steam. BWRs are largely the same as PWRs, except that there is only one water circuit, which takes heat from the fuel and produces steam [44]. All commercial reactors in the United-States are either PWR or BWR, which have very similar fuel content.

The four Canadian reactors (including the three Ontario reactors, closest to SNO+) are PHWRs, comparatively uncommon [43]. These use natural uranium as fuel, enriching the moderator (water) instead of the fuel. Pressurised heavy water is thus used in a closed circuit, similar the PWR's inner light water circuit. The Canadian reactors are of CANDU design (Canadian Deuterium-Uranium), and benefit from continous refuelling, without having to shut a whole reactor core down [44]. The three Ontario reactors (Bruce, Darlington and Pickering) make up roughly 60% of the incoming antineutrino flux, so modelling this reactor type is important.

Fission Fractions

In this analysis, the fission fractions for all PWRs are the same, taken from the averaged values of the PWR at Daya Bay [27]. The fission fractions for PHWRs were obtained via a private communication [9], and the values for both reactor types are shown in table 4.1. All other reactor types are simply modelled as PWR, due either to their fuel composition's similarity to PWR's (such as for BWR), or to distance making such details irrelevant.

As mentioned previously, these can change over time for PWRs and similar reactors, as fuel "burns up" between refuelling periods. Information on this time

Table 4.1: Isotopic fission fractions used for different reactor types, as used in equation (4.6).

Reactor Type	^{235}U	^{238}U	^{239}Pu	^{241}Pu
PWR (and others)	0.568	0.078	0.297	0.057
PHWR	0.52	0.05	0.42	0.01
PWR Uncertainty (%)	5.62	0.70	8.10	12.21

dependence is not readily available for nuclear reactors around the world, and more importantly not in North America. Instead it is treated as a systematic uncertainty in this analysis. To quantify the magnitude of fission fraction variations, standard deviations from each are computed using data from the Daya-Bay nuclear plant [45], and are also shown in table 4.1. Due to conversion from one isotope to another, the fission fractions are correlated in particular ways [40], shown in table A.1 of the appendix. These are combined to produce a fission fraction covariance matrix for PWR reactors.

However, since PHWRs are continuously refuelled – forgoing the burn-up and shutdown to refuel cycle – their fission fraction variations are far smaller than for PWRs, though a number quantifying this is not yet available. In practice, the entire reactor IBD flux will be treated as having a single fitted normalisation for this analysis. As such, ignoring oscillation, all PHWRs will always be combined into a single PDF for this analysis, which also averages out any potential variation, particularly when integrated over longer time periods. All the PWR-modelled reactors will similarly be combined into one single PDF, averaging out this variation. Therefore, the contribution of fission fraction variation to an overall reactor IBD flux normalisation uncertainty is ignored here. Any future analysis with more separately floating PDFs or high time granularity may have to account for this though.

Isotope Spectra

Until recently, semi-empirical models were used for all four isotope spectra, as described in the ILL method previously. The Huber model [46] was used for ^{235}U, ^{239}Pu and ^{241}Pu, while the Mueller model [42] was used for ^{238}U. These both provide the number of $\overline{\nu}_e$ emitted per fission, via a formula of the form

$$n(E_\nu) = \exp\left(\sum_{n=0}^{5} a_n E_\nu^n\right), \tag{4.7}$$

where a_n are fit parameters from each paper, listed in table 4.2.

However, measured reactor antineutrino fluxes by short-baseline detectors consistently find a deficit of around 6% compared to predicted models. This is known as the reactor antineutrino anomaly (RAA) [47, 45, 48], for which no explanation has been settled on yet. In addition, an excess of about 10% has been observed between around 4.5 and 6 MeV, dubbed the "5 MeV bump" [45], with no consensus on its origin[1] (see figure 4.2 below). Explaining these is not attempted in this work.

---

[1] This is not to be confused with the "low energy excess" observed by the MiniBooNE and LSND experiments. These are in the 100's of MeV range, and produced by neutrino beams [49].

Table 4.2: Isotope spectrum parameters used in equation (4.7), with a_n in units of [MeV^{-n}] [46, 42].

Isotope	a_0	a_1	a_2	a_3	a_4	a_5
^{235}U	4.367	-4.577	2.100	-0.5294	0.06186	-0.002777
^{238}U	0.4833	0.1927	-0.1283	-0.006762	0.002233	-0.0001536
^{239}Pu	4.757	-5.392	2.563	-0.6596	0.07820	-0.003536
^{241}Pu	2.990	-2.882	1.278	-0.3343	0.03905	-0.001754

Instead they are simply accounted for by taking advantage of recent measurements of short baseline reactor $\overline{\nu}_e$ IBD fluxes at the Daya Bay [50] and PROSPECT collaborations, which enabled a joint fit to be produced, with unfolded $\overline{\nu}_e$ spectra from the two primary isotope contributors [51]: ^{235}U and ^{239}Pu.

More specifically, these are the "emitted IBD spectra" for each isotope: the expected number of IBD events one would detect at the source, per fission per unit energy, with perfect detection efficiency. It also provides the smearing matrix A from the unfolding process, that one would need in order to compare other antineutrino spectra to this one, via

$$n_{\text{IBD}}\left(E_i\right) = \sum_j A_{ij} n\left(E_j\right) \sigma\left(E_j\right), \tag{4.8}$$

where i and j label the antineutrino energy bins of the provided spectra.

Now, RAT simulates IBDs by taking in the antineutrino spectra $n(E_\nu)$, so one might want to compute this by unsmearing the IBD spectrum. However, A inverts poorly since it is almost singular. So instead, the "smeared antineutrino spectrum" is computed

$$\tilde{n}\left(E_i\right) \equiv \frac{n_{\text{IBD}}\left(E_i\right)}{\sigma\left(E_i\right)} = \sum_j A_{ij} n\left(E_j\right) \frac{\sigma\left(E_j\right)}{\sigma\left(E_i\right)}, \tag{4.9}$$

so that when the IBD cross-section is added back in by RAT, the IBD spectrum is recovered – including the smearing effect. The smearing is also quite small ($A \approx \mathbf{1}$), so its impact is minimal either way. Figure 4.2 shows a comparison of these new fit spectra with the Huber model. The small excess around 5-6 MeV is clearly visible, and thus folded into the analysis.

The joint fit paper provides covariance matrices for all the bin values in the ^{235}U and ^{239}Pu spectra, including between both spectra, since they were unfolded from the same data [51]. The Mueller model paper states uncertainties and a correlation matrix for the ^{238}U fit parameters [42], which are shown in the appendix, in table A.2. From these, a covariance matrix for the ^{238}U spectrum can be computed, using the same bins as the fit ^{235}U and ^{239}Pu spectra. Lastly, the Huber model paper does not provide uncertainties for the ^{241}Pu parameters, due to high correlation making them unsuitable for error propagation [46]. Being the smallest contributor to both PWR ($\sim 6\%$) and PHWR ($\sim 1\%$) neutrino production, the uncertainty contribution from ^{241}Pu is simply ignored, assumed to be negligible.

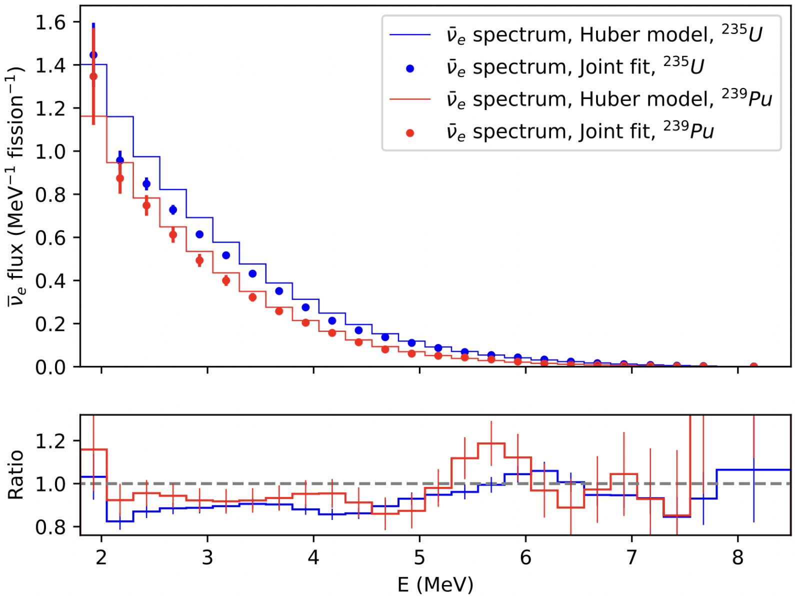

Figure 4.2: Emitted antineutrino spectra for ^{235}U and ^{239}Pu, comparing Huber model to spectrum obtained from Daya Bay and PROSPECT collaborations' joint fit [51]

Energy Emission per Fission

The average energy emitted per fission is calculated using the reactor's fission fractions, and the energy emitted per fission for each isotope from figure 4.3, as described in equation (4.5). For both reactor types considered, this amounts to

$$\begin{aligned} \epsilon_{\text{PWR}} &= (205.92 \pm 0.03)\,\text{MeV} \cdot \text{fission}^{-1}, \\ \epsilon_{\text{PHWR}} &= (206.34 \pm 0.03)\,\text{MeV} \cdot \text{fission}^{-1}, \end{aligned} \tag{4.10}$$

where the uncertainty from the fission fractions was considered negligible, as explained previously.

Table 4.3: Energy emitted per fission for each fissile isotope [41].

Fissile Isotope i	^{235}U	^{238}U	^{239}Pu	^{241}Pu
ϵ_i [MeV]	202.36 ± 0.26	205.99 ± 0.52	211.12 ± 0.34	214.26 ± 0.33

Uncertainty

Combining all of these together, one can compute the emitted PWR and PHWR spectra, per energy released. This also allows uncertainties to be properly propagated and studied. To this end, the emitted IBD spectra of both are computed, since the ^{235}U and ^{239}Pu spectra and their uncertainties are provided in this form. For the other isotopes, the $\bar{\nu}_e$ spectrum is multiplied by the IBD cross-section, though no uncertainty is yet added from the latter. This cross-section will be covered in more detail later. The spectra are shown in figure 4.3, using uncertainties, correlations, and covariance matrices as described in each sub-section above.

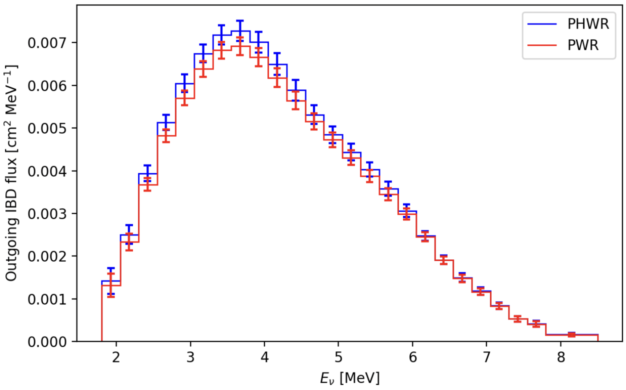

(a) Emission Spectra, with bin uncertainties.

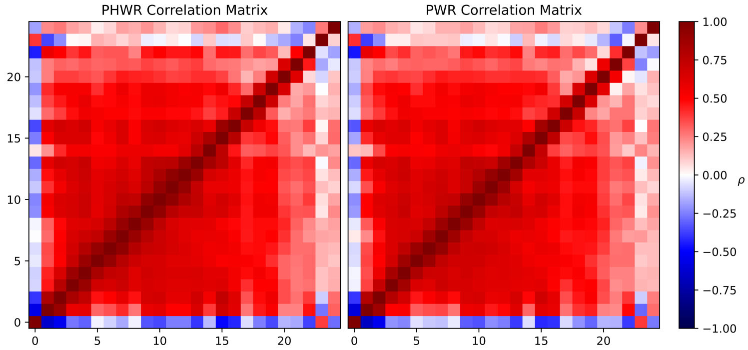

(b) Energy bin correlation matrices.

Figure 4.3: Emitted antineutrino IBD spectra for PWR and PHWR cores, and associated uncertainties.

It is clear from the correlation matrix in figure 4.3b that the energy bin uncertainties are highly (and positively!) correlated with each-other – aside from a few edge bins. Therefore, a simple overall scaling uncertainty for the whole spectrum can be used, ignoring systematic uncertainties in the shape. Integrating both spectra, the normalisation uncertainties for PHWR and PWR emission spectra are 3.1% and 2.8% respectively. This almost entirely arises from the uncertainty in the unfolded isotope spectra; the energy emitted per fission contributes negligibly to the uncertainty.

4.3.3 Power Scaling

These emitted spectra per unit energy released must be multiplied by the thermal power output of each core to obtain the true antineutrino emission spectrum: the number of emitted antineutrinos per unit time. The designed thermal power output W_{des} of all nuclear reactor cores in the world can be obtained from the IAEA yearly Operating Experience reports [52, 53]. For example, all 8 of Bruce's reactor cores are designed to output around 2500 to 2800 MW of thermal power. However, the power output of reactors are not constant over time, they vary depending on electrical grid demand, core repairs or fuel replacements. Sometimes the output even exceeds the designed amount, generally no more than by a few percent.

Scaling factors $S(t)$ must therefore be computed for each core. The IAEA reports previously mentioned also provide loading factors (LF) for every reactor core in the world, on a monthly basis. The LF of a core is defined as the ratio of its average electrical power output for that month to its reference electrical power output. This reference electrical power output is the one produced at the designed thermal power output. Therefore, if one assumes the electrical and thermal powers are proportional, this LF is also the ratio of the true monthly to designed thermal powers, and so is exactly the desired scaling factor. The time-dependent thermal power output is thus

$$W(t) = S(t)W_{\text{des}}, \tag{4.11}$$

on a monthly basis. In practice, over 99% of the incoming antineutrino flux comes from North America (US and Canada), so only these are given monthly time-dependence. Any cores outside of North America are given yearly time-dependence, since only major changes such as reactors being built or turned off permanently have any potential to impact the flux in a non-negligible way.

The three closest reactors make up around 60% of the incoming flux, and so could benefit from more fine time granularity. Fortunately, these are all in the province on Ontario, which provides hourly electrical outputs of all power plants via the Independent Electricity System Operator (IESO) [54]. The capability of these are also provided on an hourly basis. The units of the output and capability are not fully known, and the capability can vary a little over time – particularly when cores go offline or online. Therefore, the maximum capability in 2020 for each core was chosen as its reference designed output. The ratio of the hourly output and this designed output is then computed, and averaged over a day, since any more granularity than this is unnecessary. This ratio is compared to the monthly LF values from the IAEA, such as the example shown in figure 4.4. They were found to agree very well in all cases, with deviations in the average monthly values never exceeding 2%, and being around 0.2% on average. This average 0.2% deviation is taken as an additional systematic uncertainty in the reactor antineutrino flux normalisation.

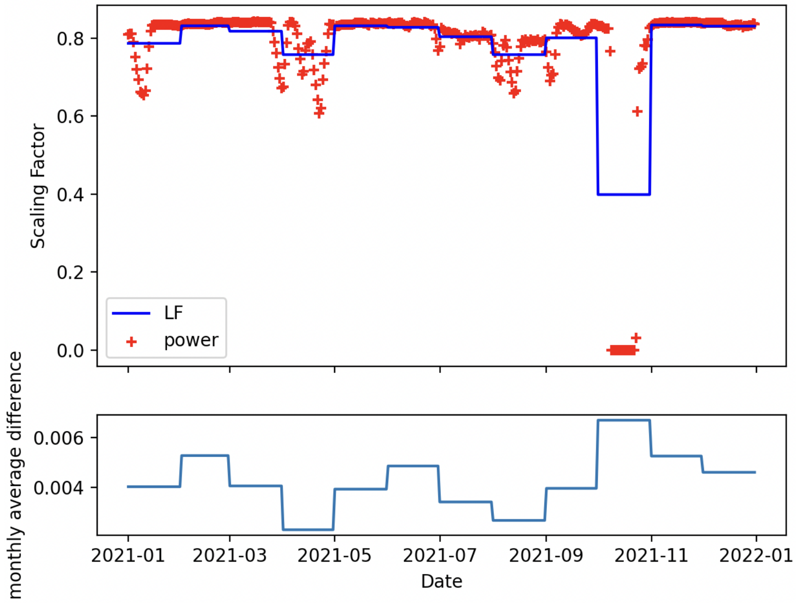

Figure 4.4: Comparison of reactor scaling factors computed using Ontario (IESO) power output data (daily, red) with those computed using IAEA loading factors (monthly, blue), for Pickering core 1.

Lastly, a note on how this scaling is implemented in RAT: 120% of the designed flux of all reactor cores is used in MC simulations, to account for reactors sometimes outputting more than their designed power. The scaling factors are then applied in offline analyses, so that they can easily be modified. The effective scaling factors used are thus related the true scaling factors such as the LF by a factor of 1.2: $S_{\text{RAT}}(t) = S(t)/1.2$. This effective scaling $S_{\text{RAT}}(t)$ is what is shown in figure 4.4. Nevertheless, this detail is irrelevant to how the analysis works, and will not be mentioned hereafter.

4.4 Event Detection

4.4.1 IBD Cross-Section

The oscillated antineutrino spectrum incident on the detector induces an IBD spectrum - what is actually measured, ignoring the detector response - via a convolution with the IBD cross-section. Only the IBD cross-section on hydrogen is considered here, since the threshold energy for larger nuclei is too high for most of the incoming antineutrinos; it would require converting any other reasonably abundant element to a less stable isotope, or breaking up the nucleus entirely. Hydrogen is also the most abundant element in the liquid scintillator. Fortunately, this makes the cross-section calculation much simpler in terms of the nuclear form factors. A historical review of this calculation is given in Ref. [55], while Ref. [56] provides a summary of various approximations, and their associated uncertainties. These approximations are generally valid for either low ($\lesssim$ 60MeV) or high (GeV range) energies, though the latter paper also produces a more accurate expression valid across the whole energy range. For the case of reactor antineutrinos, the low energy approximation from Ref. [57] is more than accurate enough, as will be seen in the error analysis below.

A walk-through of the calculation of a low-energy approximation is presented first below, to gain a more qualitative understanding of the process, as well as how the neutron lifetime enters the calculation. An expansion of the positron energy in terms of nucleon masses is then provided, laying the groundwork for a description of the more accurate differential cross-section by Ref. [57]. The author's own calculation of the total cross-section then follows. This subsection is finished off by a discussion of the quantities entering the calculation and their respective uncertainties.

Low Energy Approximation

Fermi Theory Looking at the IBD's Feynman diagram in figure 4.1, one can see that it involves the leptonic Weak nuclear interaction term from the Lagrangian 2.3, as well as the similar term for up and down quarks,

$$\mathcal{L}_{\text{int}} = i\frac{g_W}{2\sqrt{2}}\left[\bar{\nu}_e \not{W}\left(1-\gamma^5\right)e + V_{ud}^*\bar{d}\not{W}\left(1-\gamma^5\right)u + \text{h.c.}\right]. \tag{4.12}$$

V is the Cabibbo–Kobayashi–Maskawa (CKM) matrix governing the flavour mixing between quarks – in the same way the PMNS matrix does for neutrinos and charged leptons – so that V_{ud} describes the mixing between up and down quarks. The reason the CKM matrix is present here, but the PMNS matrix is not, is that quark flavours are all defined according to their respective mass states while neutrino flavour states are chosen to match their respective charged leptons. Essentially the PMNS matrix was chosen to appear in the neutrino's mass terms instead of here. One can find the associated second-order transition amplitude via effective 4-fermion Fermi theory, due to the low energies involved ($E_\nu \ll M_W$), so that

$$\mathcal{A}_{p,\bar{\nu}_e \to n,e^+} = -iz\frac{G_F}{\sqrt{2}}V_{ud}^* \left[\bar{v}^{s_{\nu_e}}(p_{\nu_e})\gamma^\mu \left(1-\gamma^5\right) v^{s_e}(p_e)\right] \left[\bar{u}^{s_n}(p_n)\gamma_\mu \left(f - g\gamma^5\right) u^{s_p}(p_p)\right], \tag{4.13}$$

where z includes all energy-independent inner radiative corrections, G_F is Fermi's constant, and $u^s(p)$ and $v^s(p)$ are the spinor and anti-spinor fields introduced in (2.9). The proton and neutron 4-momenta are used here to account for the fact that the up and down quarks are not free, but are merely the valence quarks in their respective hadrons [4]. f and g are the effective vector and axial coupling constants respectively, introduced to quantify the incalculable effects of the strong force and the hadrons' internal structures. These only work in the low-energy limit, such as for neutron decay or approximately for low-energy IBDs. Outside this limit, these must be replaced by more complex terms, involving other quantities such as the nucleon anomalous magnetic moment, for the most accurate calculations [4]. This is an example of the so-called V-A (vector-axial) structure used to described weak interactions. The weak leptonic interaction between the antineutrino and the positron has the same structure, but with $f = g = 1$ identically, due to being an interaction of purely left-handed Weyl spinors not part of any bound states, as previously discussed in section 2.3.1.

To lowest order in the nucleon mass, one can show that this yields the differential cross-section

$$\left(\frac{d\sigma}{d\cos\theta}\right)^{(0)} = \frac{\sigma_0}{2}\frac{E_\nu^2}{v_e}\left[(f+g)^2 + \left(f^2 - g^2\right)v_e\cos\theta\right], \tag{4.14}$$

where $v_e = |\vec{p}_e|/E_e$, θ is the angle between the incoming antineutrino and the outgoing positron,

$$\sigma_0 = z^2\frac{G_F^2\cos^2\theta_C}{\pi}, \tag{4.15}$$

and $|V_{ud}|^2 = \cos^2\theta_C$ is the Cabibbo angle.

Using Neutron Decays The constants in σ_0 have been measured primarily using neutron decays, since the interaction cross-section is identical, aside from differing kinematics. This is called crossing symmetry, and can be seen either by comparing their Feynmann diagrams, or by looking at what they represent: transition amplitudes. The neutron decay transition amplitude is almost identical to 4.13, with just a few tweaks

$$\mathcal{A}_{n \to p,e,\nu} = -iz\frac{G_F}{\sqrt{2}}V_{ud}^* \left[\bar{u}^{s_e}(p_e)\gamma^\mu \left(1-\gamma^5\right) v^{s_\nu}(p_\nu)\right] \left[\bar{u}^{s_p}(p_p)\gamma_\mu \left(f - g\gamma^5\right) u^{s_n}(p_n)\right], \tag{4.16}$$

which one can similarly show yields a formula for the neutron lifetime of the form [4]

$$\tau_n = \frac{2\pi^3\left(f^2+g^2\right)^{-1}}{G_F^2\cos^2\theta_C m_e^5 f_{\text{p.s.}}^R z^2}, \tag{4.17}$$

where $f^R_{\text{p.s.}}$ is the dimensionless phase-space integral, taking into account Coulomb, weak magnetism, recoil, and outer – but not inner – radiative corrections [58].

One can therefore write the σ_0 constant directly in terms of the measured neutron lifetime, as

$$\sigma_0 = \frac{2\pi^2 \left(f^2 + g^2\right)^{-1}}{\tau_n m_e^5 f^R_{\text{p.s.}}}, \tag{4.18}$$

absorbing all the inner radiative corrections into a quantity which is affected by them in the same way.

Antineutrino and Positron Energies

Before discussing the full cross-section formula, some kinematic quantities are considered. In the lab frame, the incoming electron antineutrino interacts with a proton assumed to be at rest, producing a neutron and a positron. With the usual approximation that the antineutrino is massless, one can show that

$$E_\nu = \frac{1}{2M_p}\left[2E_eE_n + M_n^2 - M_p^2 + m_e^2 - 2\vec{p}_e \cdot \vec{p}_n\right], \tag{4.19}$$

from which one can determine the minimum antineutrino energy needed to trigger the decay,

$$E_\nu^{\text{min}} = \frac{\left(M_n + m_e\right)^2 - M_p^2}{2M_p} \approx 1.806 \text{ MeV}. \tag{4.20}$$

This is approximately the difference between the neutron and proton masses, as one might expect. Now, rearranging (4.19) to a formula for E_e, and using the angle between the incoming antineutrino and outgoing positron momenta θ, one finds

$$E_e = \frac{1}{E_n}\left[M_n\left(E_\nu - \Delta\right) - E_eE_\nu\left(1 - v_e\text{cos}\theta\right) + E_\nu\left(E_e - \Delta\right) - E_e^2 + \frac{1}{2}\left(\Delta^2 + m_e^2\right)\right], \tag{4.21}$$

where $\Delta = M_p - M_n$. Expanding $\frac{E_\nu}{E_n}$ in orders of $\frac{E_\nu}{M_n}$, one can find the zeroth and first order approximations for the positron energy

$$E_e^{(0)} = E_\nu - \Delta, \tag{4.22}$$

and

$$E_e^{(1)} = E_e^{(0)}\left[1 - \frac{E_\nu}{M_n}\left(1 - v_e^{(0)}\text{cos}\theta\right)\right] - \frac{y^2}{M_n}, \tag{4.23}$$

where $y^2 = \frac{1}{2}\left(\Delta^2 - m_e^2\right)$. These are referenced in the cross-section calculation below, and the first order formula is used in simulations. Note however, that this can occasionally allow for energies below m_e. Any such positrons are instead simulated with an energy of m_e to maintain physicality.

The Differential Cross-Section

As previously mentioned, the SNO+ experiment uses the differential cross-section, to first order in E_ν/M, where M is the nucleon mass, from [57]. This is given by

$$\left(\frac{d\sigma}{d\text{cos}\theta}\right)^{(1)} = \frac{\sigma_0}{2}\left[\left(f^2 + 3g^2 + \left(f^2 - g^2\right)v_e^{(1)}\text{cos}\theta\right)E_e^{(1)}p_e^{(1)} - \frac{\Gamma}{M}E_e^{(0)}p_e^{(0)}\right], \tag{4.24}$$

where

$$\begin{aligned}\Gamma = 2\left(f+f_2\right)g&\left[\left(2E_e^{(0)}+\Delta\right)\left(1-v_e^{(0)}\cos\theta\right)-\frac{m_e^2}{E_e^{(0)}}\right]\\
&+\left(f^2+g^2\right)\left[\Delta\left(1+v_e^{(0)}\cos\theta\right)+\frac{m_e^2}{E_e^{(0)}}\right]\\
&+\left(f^2+3g^2\right)\left[\left(E_e^{(0)}+\Delta\right)\left(1-\frac{1}{v_e^{(0)}}\cos\theta\right)-\Delta\right]\\
&+\left(f^2-g^2\right)\left[\left(E_e^{(0)}+\Delta\right)\left(1-\frac{1}{v_e^{(0)}}\cos\theta\right)-\Delta\right]v_e^{(0)}\cos\theta,\end{aligned}\tag{4.25}$$

with all the same variable definitions as previously, and $f_2 = \bar{\mu}_p - \bar{\mu}_n$ is the anomalous nucleon isovector magnetic moment. This approximation works extremely well at the energies of interest – higher order corrections contribute negligible uncertainty compared to that introduced by the σ_0 normalization constant, described below. It is also computationally faster than the more precise formula in [56], hence its use in simulations at SNO+. As for higher order Feynman diagram contributions, these would be suppressed by an additional factor of G_F^2, contributing even less uncertainty.

The Total Cross-Section

To compute the total expected IBD flux, the differential cross-section must be integrated over all $\cos\theta \in \{-1, 1\}$, to obtain the total cross-section. Note that $\Gamma = \Gamma(\cos\theta)$ and $E_e^{(1)} = E_e^{(1)}(\cos\theta)$ have $\cos\theta$-dependence, so $p_e^{(1)} = p_e^{(1)}(\cos\theta)$ and $v_e^{(1)} = v_e^{(1)}(\cos\theta)$ do too, while the zeroth order quantities do not. Further notice that any odd power of $\cos\theta$ will integrate to zero, since it is antisymmetric about $\cos\theta = 0$. Using this and the following facts

$$v_e^{(1)}E_e^{(1)}p_e^{(1)} = E_e^{(1)2} - m_e^2, \qquad \frac{d\cos\theta}{dE_e^{(1)}} = \frac{M}{p_e^{(0)}E_\nu}, \tag{4.26}$$

as well as only keeping terms up to first order in $\frac{E_\nu}{M}$, one can show that

$$\sigma^{(1)} = \sigma^{(0)}\left(1+\frac{\Gamma}{M}\right), \tag{4.27}$$

where

$$\sigma^{(0)} = \frac{2\pi^2}{m_e^5 f_{\text{p.s.}}^R \tau_n}E_e^{(0)}p_e^{(0)}, \tag{4.28}$$

and

$$\begin{aligned}\Gamma =&\frac{2}{\left(f^2+3g^2\right)}\left[g\left(f+g+f_2\right)\left(\Delta-2E_\nu\right)+\frac{m_e^2}{E_e^{(0)}}\left(gf_2-\frac{1}{2}\left(f-g\right)^2\right)\right]\\
&-\left(1+\frac{1}{v_e^{(0)2}}\right)\left(E_\nu-\frac{y^2}{E_e^{(0)}}\right).\end{aligned}\tag{4.29}$$

Indeed, even the original differential cross-section (4.24) technically had some higher than first order terms from $E_e^{(1)}p_e^{(1)}$. In the limit $\frac{E_\nu}{M} \to 0$ this reduces to the zeroth order total cross-section $\sigma^{(0)}$, matching up with the previous calculation in (4.14), if the neutron is at rest ($|\vec{p}_n|/|\vec{p}_e| \to 0$), which is also implied by the low energy limit. Lastly, one should bear in mind that the differential cross-section must still be used to simulate positrons with the correct energy distribution, since $E_e^{(1)}$ depends on $\cos\theta$.

Constants and Uncertainty

Recall the normalisation constant can be defined either according to G_F, $\cos\theta_C$ and the inner radiative corrections $z^2 = 1 - \Delta^R_{\text{inner}}$ (4.15), or according to the neutron life time τ_n and its associated phase space integral $f^R_{\text{p.s.}}$ (4.18). To choose which definition to use, one must look at how the quantities are measured. The Cabibbo angle θ_C is taken from measurements of V_{us} via Kaon semi-leptonic decay (K_{l3}), and from nuclear and neutron decays [56]. These two measurements only agree at the 2.2σ level, increasing the uncertainty in θ_C. Meanwhile, the ratio g/f and Δ^R_{inner} are also measured from neutron decay experiments, and have to take $\cos\theta_C$ and τ_n as known parameters in their calculations. For examples on measuring Δ^R_{inner}, see [59] or [60], of which the former also uses results from β decays of larger atoms (there appears to be quite a spread of values between different atoms). It was thus decided to simply use the neutron lifetime measurement directly, since it only involves a free neutron and proton just like the IBDs considered here, and cuts away the complex correlations and tensions between some of these measurements, particularly as some are updated over time. The neutron decay time is [61]

$$\tau_n = (879.4 \pm 0.6)\, s. \tag{4.30}$$

Now turning to the other constants involved, $f = 1$ by definition (ignoring isospin-breaking corrections, which are negligible [56]). As mentioned above, the value of g (technically the ratio g/f) comes from neutron decay measurements, yielding [62]

$$g = 1.27641 \pm 0.00056. \tag{4.31}$$

as does the phase space factor, so that [58]

$$f^R_{\text{p.s.}} = 1.71517 \pm 0.00009. \tag{4.32}$$

The anomalous nucleon isovector magnetic moment f_2 is defined according to the the anomalous magnetic moments of the proton and neutron, in units of the nucleon magentic moment. From [6], the proton and neutron magnetic moments are $\mu_p = 2.7928473446 \pm 0.0000000008$ and $\mu_n = -1.9130427 \pm 0.0000005$, whereas the expected Dirac particle magnetic moments are 1 for the proton and 0 for the neutron. Subtracting these from the measured magnetic moments gives the anomalous magnetic moments, and the difference between these defines $f_2 = \bar{\mu}_p - \bar{\mu}_n$, so that

$$f_2 = 3.7058900 \pm 0.0000005. \tag{4.33}$$

Lastly, the electron mass and proton masses are also taken from [6]

$$m_e = (0.51099895000 \pm 0.00000000015)\,\text{MeV}, \tag{4.34}$$

$$M_p = (938.27208816 \pm 0.00000029)\,\text{MeV}, \tag{4.35}$$

and the proton mass is used as the nucleon mass, since the difference between E_ν/M and E_ν/M_p is of order E_ν/M^2.

Propagating these uncertainties, one finds an uncertainty in the zeroth order total cross-section of 0.07%, primarily from the neutron lifetime. The extra quantities involved in the differential cross-section do not greatly change this, and neither do terms of higher order in E_ν/M, simply due to the energies involved. The authors of [56] consider the uncertainties from both definitions of σ_0, and provide a conservative estimate of the uncertainty, which is adopted for this work

$$\sigma_{\sigma_{\mathrm{IBD}}} = 0.4\%, \tag{4.36}$$

and will turn out to be a sub-dominant source of uncertainty either way.

4.4.2 The Target Protons

Scintillator Density

The number of protons in a chosen SNO+ detector volume is directly proportional to the scintillator density, which itself depends on temperature and pressure. It must therefore be carefully determined to eliminate it as a primary source of uncertainty. Such a measurement was performed by B. Tam in the summer of 2022, in Snolab's underground laboratory, for various scintillator temperatures, as shown in figure 4.5. The air pressure during these measurements was 0.12 MPa, and variations in this (±5%) are expected to have a negligible impact. A linear fit was applied to this data, providing the following temperature dependence of the liquid scintillator density,

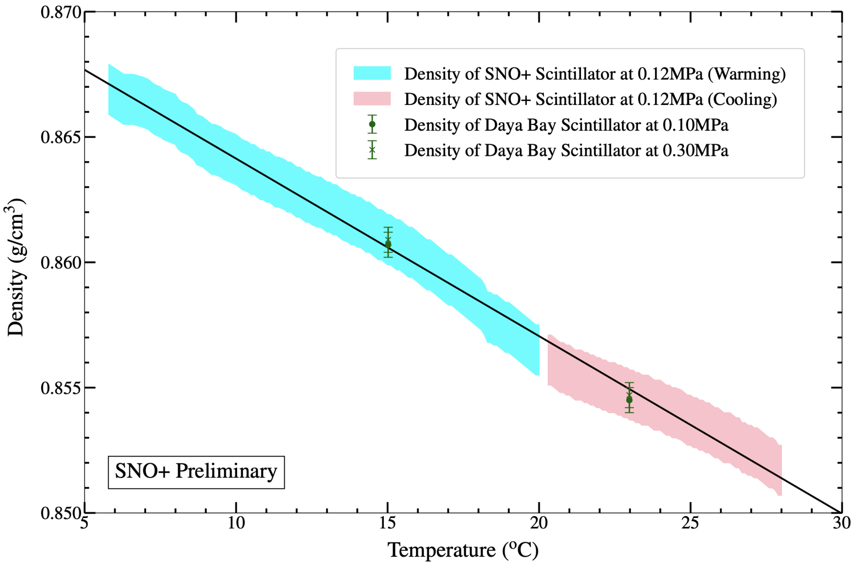

Figure 4.5: Density measurements of SNO+ sinctillator samples taken with the Anton Paar DMA 35 portable density and concentration meter, at regular 0.1°C intervals in the 5.8°C to 20°C and 20.3°C to 28°C ranges. Uncertainties in these measurements are displayed as blue and red bands. A linear fit is overlain, along with published Daya Bay scintillator density measurements for comparison [63, 64].

$$\rho(T)[\mathrm{g/cm^3}] = -0.00070831 \times T[^\circ\mathrm{C}] + 0.87122, \tag{4.37}$$

where T is the temperature [63]. This fit appears in excellent agreement with measurements of Daya Bay's similar LAB based scintillator, even at differing pressures. The linear fit had a coefficient of determination of $R^2 = 0.9833$, indicating an excellent fit, while the datapoints had errors of around 0.2% or less. The fit itself thus contributes negligibly to uncertainty in the number of protons.

The temperature inside the AV cannot be measured directly, so the temperature of the external water is measured instead, with sensors mounted onto the PSUP at various positions. These display temperatures ranging from 12°C at the bottom of the PSUP, to 18°C at the top, with an average of around 15°C towards the middle. These temperatures are quite stable, very rarely exceeding variations of ±1°C over time. The scintillator density at 15°C is used in this analysis, with a conservative uncertainty of ±3°C, so that

$$\rho = (0.8606 \pm 0.0021)\,\mathrm{g/cm^3}, \tag{4.38}$$

which is a 0.24% uncertainty. Also, it must be noted that RAT uses the density at 12°C (0.8627 g/cm^3), so a scaling of 0.24% must be applied to any simulation results that depend on this.

Hydrogen Fraction

The obtain the proton density from this, one must compute the hydrogen mass fraction of the liquid scintillator cocktail. As discussed in section 3.2, the SNO+ scintillator is made up primarily of LAB, with 2.2g/L of PPO, and also $\mathcal{O}(1)$ mg/L or less of bis-MSB (1,4-bis(2-methylstyryl)benzene) and BHT (butylated hydroxytoluene), depending on the time. Fortunately, the quantity of bis-MSB and BHT renders them completely negligible in this calculation.

The 2.2g/L of PPO translates to a mass fraction of $\alpha_{\mathrm{PPO}} = 0.253\%$ of the scintillator cocktail, while PPO has a hydrogen mass fraction of $f_{\mathrm{PPO}} = 5.01\%$. The SNO+ LAB is made up of various carbon chains, listed in table 4.4. Bringing this all together, the scintillator's hydrogen mass fraction is

$$f = 12.17\%, \tag{4.39}$$

and although an uncertainty is difficult to estimate, it is almost certainly negligible. Due to the similarity in the carbon chains' hydrogen mass fractions f_i (especially the two dominant ones), even an uncertainty in all their mass fractions α_i of 100%, far beyond the manufacturer's limits, would produce only a 0.45% uncertainty in f. This is while enforcing the normalisation condition $\sum_i \alpha_i = 1$, which must hold by definition. Likewise, an uncertainty in α_{PPO} of 100%, once again very far beyond any reasonable uncertainty in the mass of loaded PPO, only results in an uncertainty of 0.27%. A dedicated measurement of the hydrogen and carbon fractions of the scintillator has not been performed however, but would be advantageous to confirm this calculation and assign a definite uncertainty to it.

Therefore, combining this hydrogen mass fraction with the previously computed scintillator density, and a hydrogen atomic mass of 1.0078 amu [66], the proton density in SNO+ liquid scintillator is

$$n_H = 6.251 \times 10^{28}\mathrm{m}^{-3} \pm 0.3\%, \tag{4.40}$$

Table 4.4: Mass fractions of carbon chains in SNO+ LAB (α_i) and the manufacturer's stated limits on these [65], along with the hydrogen mass fractions for each chain (f_i).

Carbon chain	α_i (%)	f_i (%)	Manufacturer limits
$C_{15}H_{24}$	0.9	11.85	$< 3\%$
$C_{16}H_{26}$	18	12.01	$< 25\%$
$C_{17}H_{28}$	45.5	12.15	$(C_{16}H_{26} + C_{17}H_{28}) > 50\%$
$C_{18}H_{30}$	30.4	12.28	$> 25\%$
$C_{19}H_{32}$	5.2	12.39	$< 15\%$

where a conservative 0.3% uncertainty is assigned, which is negligible. It turns out that only events inside a volume defined by a 5.7m sphere centered on the centre of the detector are considered, to reduce backgrounds. The reason for this is explained in chapter 6. The total number of target protons is thus

$$N_H = 4.849 \times 10^{31} \pm 0.3\%. \tag{4.41}$$

The uncertainty of events' reconstructed positions is roughly 10 mm [67], though it is assumed that this affects both real data and MC simulation in the same way. Therefore, this is accounted for in the tagging efficiency calculation of chapter 6. The volume itself thus does not contribute any systematic uncertainty, though a more careful consideration of this last point should be carried out in the future.

4.5 Summary

Putting all this together, one can compute the expected reactor IBD spcectrum at SNO+, ignoring for now the impact of detection efficiency and neutrino oscillation. This is shown in figure 4.6 – in neutrino energy, since the conversion to the detected prompt energy is performed event-by-event in RAT's MC simulations, incorporating reconstruction and detection efficiency effects. The expected prompt energy spectrum will be shown in the Backgrounds chapter, in the context of cuts and detection efficiencies. However, to first order, these two energies are related by $E = E_\nu - 1.29[\text{MeV}]$, as described previously. Nevertheless, this integrates to an expectation of 58.9 reactor IBDs for the run period under consideration, where the flux contributions from PHWRs and PWRs are split roughly 56% to 44%, due to their real-time power scalings. Finally, note that some real-time scaling information was not available for some reactors at the time of this analysis, so the overall normalisation is liable to change slightly as that becomes available.

The only non-negligible uncertainty on this is an overall normalisation uncertainty, as discussed previously. The combination of all the contributions to this is shown in table 4.5. The PHWR and PWR uncertainties are quite similar, around 3%. Aside from the 0.2% power scaling uncertainty, the uncertainties are all fully correlated for all reactor core sources. Lastly, the effect of the conversion from antineutrino energy to positron energy is effectively an energy shift and an order $\mathcal{O}(1/M)$ smearing. This, in combination with the residual energy scaling and smearing from reconstruction have a negligible impact on the normalisation, and so do

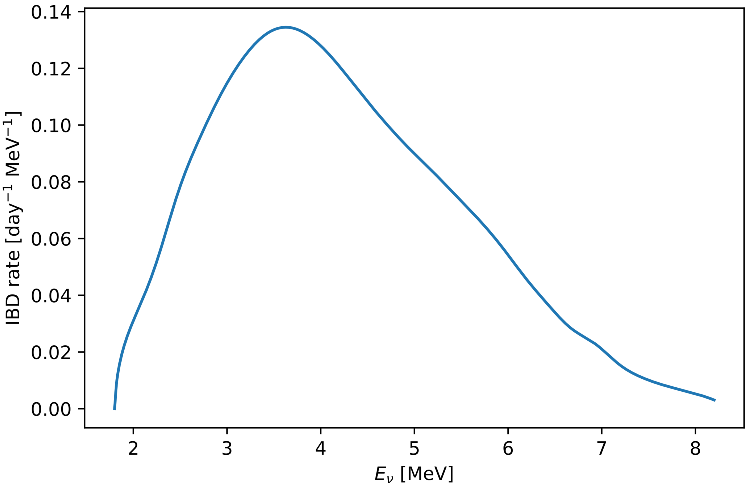

Figure 4.6: Expected reactor IBD spectrum at SNO+, in antineutrino energy, ignoring neutrino oscillation, detector and tagging efficiencies, during the 134.5 days of this analysis.

not contribute to its uncertainty. Therefore, a single 3% normalisation uncertainty for all reactor IBDs is employed in this analysis

Table 4.5: Reactor IBD normalisation uncertainties.

Source	Uncertainty [%]	
	PHWR	PWR
Isotope spectra	3.1	2.8
Power scaling	0.2	
IBD cross-section	0.4	
Number of protons	0.3	
Total	3.1	2.9

Chapter 5

Neutrino Oscillation Phenomenology

> “Home is behind, the world ahead,
> And there are many paths to tread
> Through shadows to the edge of night,
> Until the stars are all alight.”
>
> *The Lord of the Rings*, J.R.R. Tolkien

5.1 Introduction

The reactor IBD flux and spectrum detected at SNO+ is greatly affected by neutrino oscillation, which is why oscillation can be measured using them. This chapter delves into the details of the approximations and phenomenology of neutrino oscillation, and the history of their measurements, particularly as they pertain to the setup at SNO+. The general $\overline{\nu}_e$ survival probability in vacuum formula (2.24) was shown already in chapter 2, as an introduction to the topic. Written out explicitly here in terms of the usual PMNS matrix parametrisation,

$$\begin{aligned} P_{\overline{\nu}_e\rightarrow\overline{\nu}_e}(L,E) = 1 - 4\Bigg[&\left(s_{12}c_{12}\right)^2 c_{13}^4 \sin^2\left(\Delta m_{21}^2 \frac{L}{4E}\right) \\ &+ \left(s_{13}c_{12}c_{13}\right)^2 \sin^2\left(\Delta m_{31}^2 \frac{L}{4E}\right) \\ &+ \left(s_{13}s_{12}c_{13}\right)^2 \sin^2\left(\Delta m_{32}^2 \frac{L}{4E}\right)\Bigg], \end{aligned} \tag{5.1}$$

it serves as the starting point for what follows. In this chapter, E refers only to (anti)neutrino energy, while L is the baseline as usual.

The long and short baseline approximations are covered first, followed by an introduction to the impact of propagation through matter. Some of the history of neutrino oscillation measurements is then reviewed, to explain the current status of relevant parameter measurements. A novel calculation is then presented which includes the effects of (anti)neutrinos travelling through constant non-zero matter density, providing a more accurate and fast survival probability calculation for SNO+.

5.2 Baseline Approximations in Vacuum

Notice that equation (5.1) oscillates over L/E, with different wavelengths given by $4\pi/\Delta m^2_{21}$, $4\pi/\Delta m^2_{31}$ and $4\pi/\Delta m^2_{32}$. However, only two of these mass differences are independent by definition, and in practice Δm^2_{21} is far smaller than the other two: $\Delta m^2_{21} \ll \left|\Delta m^2_{31}\right| \approx \left|\Delta m^2_{32}\right|$. Therefore, these are often referred to as two mass differences: the solar and atmospheric mass differences, respectively, after how they were originally measured, so that $\Delta m^2_{SOL} \ll \Delta m^2_{ATM}$. So far this and other measurements allow for two possible mass orderings [4], shown in figure 5.1.

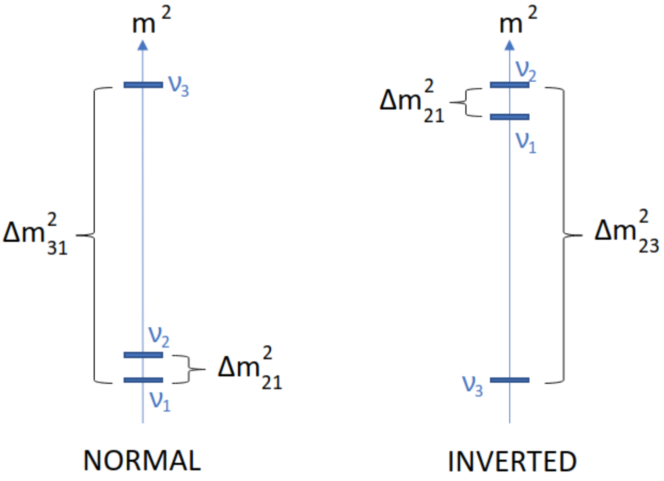

Figure 5.1: Possible mass orderings

The oscillation this produces is shown in figure 5.2a, where one long wavelength and one short wavelength are most obvious, corresponding to $4\pi/\Delta m^2_{21} \approx 33\ \mathrm{km} \cdot \mathrm{MeV}^{-1}$ and $4\pi/\Delta m^2_{31} \approx 1\ \mathrm{km} \cdot \mathrm{MeV}^{-1}$. These very different scales necessitate different experimental setups to be most sensitive: long baseline and short baseline experiments. In short baseline experiments, the detector is placed close to the source, so that the longer oscillation has no impact. In long baseline experiments, the detector is placed much further away, where the small-scale oscillation becomes averaged out and negligible by comparison.

In these regimes, one can derive the approximate behaviours from (5.1), so that for long baselines

$$P_{\overline{\nu}_e \to \overline{\nu}_e}(L, E) = 1 - \frac{1}{2}\sin^2\left(2\theta_{13}\right) - \sin^2\left(2\theta_{12}\right) c^4_{13} \sin^2\left(\Delta m^2_{21}\frac{L}{4E}\right), \tag{5.2}$$

where the Δm^2_{31} and Δm^2_{32} oscillation terms were averaged over via $\langle \sin^2\theta \rangle = \frac{1}{2}$. This shows that long baseline experiments are also most sensitive to θ_{12} out of all the mixing parameters, particularly since $\sin^2\left(2\theta_{13}\right)$ happens to be close to zero, as will be discussed shortly. One can additionally compute a short baseline approximate formula by using $\Delta m^2_{32} = \Delta m^2_{31} - \Delta m^2_{21}$, followed setting any $\sin\left(\Delta m^2_{21}\frac{L}{4E}\right) = 0$ and $\cos\left(\Delta m^2_{21}\frac{L}{4E}\right) = 1$. Recall, there are only two independent mass differences between three masses. With some trigonometric identities one can show that this becomes

$$P_{\overline{\nu}_e \to \overline{\nu}_e}(L, E) = 1 - \sin^2\left(2\theta_{13}\right) \sin^2\left(\Delta m^2_{31}\frac{L}{4E}\right), \tag{5.3}$$

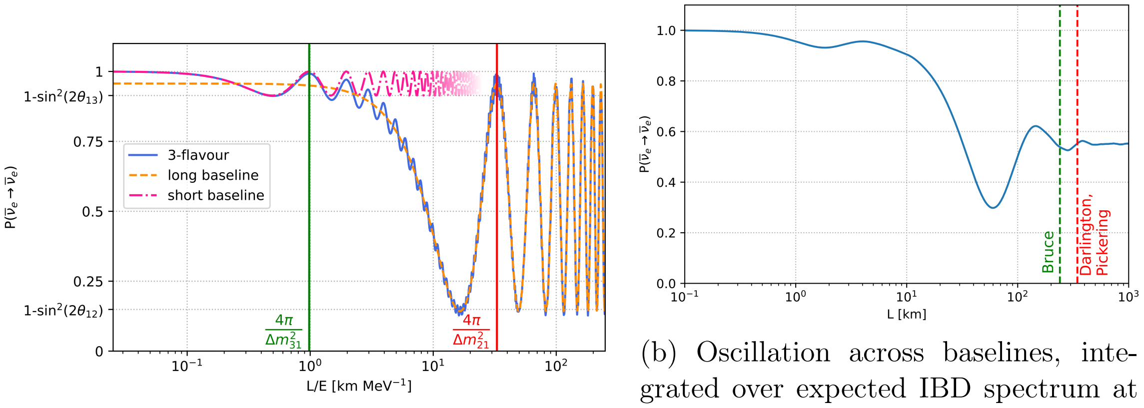

(a) Oscillation across L/E scales, with the full three-flavour calculation, and the two-flavour long and short baseline approximations.

(b) Oscillation across baselines, integrated over expected IBD spectrum at SNO+, from figure 4.6. Approximate distances to the nearest three nuclear reactors are also shown for reference.

Figure 5.2: Vacuum oscillation probabilities, using PDG 2022 values [6].

from which it is evident that short baseline experiments are instead only sensitive to the θ_{13} mixing angle.

In either of these two baseline approximations, the three-flavour oscillation has been reduced to two-flavour oscillation with new effective mixing angles. Also notice that they are completely insensitive to the mass ordering, due to the mass differences only appearing in $\sin^2()$ terms, and the previously mentioned $\left|\Delta m^2_{31}\right| \approx \left|\Delta m^2_{32}\right|$ approximation. More precise calculations with the full three-flavour formula must be used to see the small differences this produces. Alternatively, the impact of very high mass densities can be leveraged, as will be described later. Short baseline appearance probabilities (i.e. the transition from one flavour to a single other flavour) could also potentially show sensitivity to this. Regardless, the above approximations are also displayed in figure 5.2a, where one can see that as a result, the small θ_{13} governs the short baseline oscillation amplitude, while the larger θ_{12} governs the amplitude of the long baseline oscillation.

With the closest three nuclear reactors to SNO+ being around 240km and 350km away, and making up roughly 60% of the incoming antineutrino flux (centered around 3-5 MeV), the SNO+ detector is a sensitive long baseline detector. This can be seen in figure 5.2b, where the survival probability was integrated over the expected IBD spectrum from the previous chapter. The closest reactors are around the “second peak”, or dip in this case, showing sensitivity of the flux normalisation to Δm^2_{21} before the survival probability gets averaged out to roughly 0.54 beyond around 500km. This value is obtained when the long baseline formula (5.2) is further averaged out over L/E, producing

$$\langle P_{\overline{\nu}_e \to \overline{\nu}_e} \rangle = s^4_{13} + c^4_{13}\left[1 - \frac{1}{2}\sin^2\left(2\theta_{12}\right)\right] \approx 0.54. \tag{5.4}$$

This occurs with small averaging over E at large L such as shown in figure 5.2b, or when averaged over large enough E or L, as will be the case for geo-neutrinos covered in the next chapter.

The sensitivity of SNO+ goes beyond this however: since these reactors are at a critical distance form the detector, antineutrinos at different energies have drastically different survival probabilities. This means the IBD spectrum’s shape is greatly sensitive to Δm^2_{21}, far more than its overall scale. Examples of the impact of changing Δm^2_{21} on the IBD spectrum are shown in figure 5.8b. Meanwhile, θ_{12} has the larger impact on the flux scale, making SNO+ sensitive to both parameters in complementary ways: Δm^2_{21} by spectral shape, and θ_{12} by total flux.

5.3 Matter Effects

All the oscillation calculations and arguments above have assumed that neutrinos are travelling through a vacuum, when in many realistic cases they traverse a great deal of matter, either in the Earth, the Sun, or interstellar dust. From basic dimensional arguments, one can show that the cross section for incoherent scattering of neutrinos is of the order $\sigma \sim G_F^2 s$, just like the IBD cross-section computed in the previous chapter. As such, it is exceedingly rare in almost all realistic cases, requiring neutrino energies of at least 10^5 GeV inside the Earth, or neutron star densities at the MeV scale, to be non-negligible [4]. Instead, it is coherent forward elastic scattering that turns out to have a measurable impact, acting as a refractive index that modifies neutrino oscillation, as will be shown below.

5.3.1 Theoretical Motivation

Recall that in the previous derivation of neutrino oscillations, only the free Lagrangian was used to construct the free spacetime translation operator $\hat{P}_0^\mu$ (2.16). Strictly speaking, one should use the full Lagrangian, including what leads to the interaction terms, so that

$$\begin{aligned} \partial^\mu |\nu_\alpha(x)\rangle &= \hat{P}^\mu |\nu_\alpha(x)\rangle \\ &= \left(\hat{P}_0^\mu + \hat{P}_I^\mu\right) |\nu_\alpha(x)\rangle \\ &\approx \left(\hat{P}_0^\mu + \left\langle \mathcal{A}_{\nu\to\nu}\right\rangle\right) |\nu_\alpha(x)\rangle, \end{aligned} \tag{5.5}$$

where $\left\langle \mathcal{A}_{\nu\to\nu}\right\rangle$ is the average effective interaction amplitude, caused by coherent forward elastic scattering in matter. In non-exotic matter these can only be the two processes in figure 5.3, since the initial and final states are required to be identical, and only the lowest order interaction vertices need be considered. These CC and NC transition amplitudes can be computed similarly to the IBD one from equation (4.13). In the low energy limit of Fermi theory, all these vertices have the same vertex coupling, except that the CKM matrix drops out: for the CC branch because no quarks are involved, and for the NC branch because the Z boson does not lead to flavour mixing. The effective vector and axial coupling constants f and g also need to be adapted to the particular fermions in question. Aside from these caveats, one can simply replace the relevant spinors in one of the vertex terms of (4.13), so that

$$\begin{aligned} \mathcal{A}_{CC} &= -i\frac{G_F}{\sqrt{2}} \left[\bar{v}^{s_{\nu_e}}(p_{\nu_e})\gamma_\mu \left(1-\gamma^5\right) v^{s_e}(p_e)\right] \left[\bar{v}^{s_e}(p_e)\gamma^\mu \left(1-\gamma^5\right) v^{s_{\nu_e}}(p_{\nu_e})\right], \\ \mathcal{A}_{NC} &= -i\frac{G_F}{\sqrt{2}} \sum_{\alpha=e,\mu,\tau} \left[\bar{v}^{s_{\nu_\alpha}}(p_{\nu_\alpha})\gamma_\mu \left(1-\gamma^5\right) v^{s_{\nu_\alpha}}(p_{\nu_\alpha})\right] \\ &\qquad \times \sum_{\beta=e,p,n} \left[\bar{v}^{s_\beta}(p_\beta)\gamma^\mu \left(f_\beta - g_\beta\gamma^5\right) v^{s_\beta}(p_\beta)\right]. \end{aligned} \tag{5.6}$$

The NC amplitude was summed over all possible particle combinations, and the radiative correction factor was simply set to $z = 1$, since such high accuracy is not required in this case. Averaging over all the particles in the matter medium, setting their initial and final spins equal, and taking into account their statistical energy distribution in matter, one can show that the neutrinos feel an overall potential [4]

$$\langle \mathcal{A}_{\nu\to\nu}\rangle|\nu_\alpha\rangle = -iV_\alpha|\nu_\alpha\rangle, \quad V_\alpha = \pm V_{CC}\delta_{\alpha e} + V_{NC} = \sqrt{2}G_F\left(\pm N_e\delta_{\alpha e} - \frac{1}{2}N_n\right), \tag{5.7}$$

where N_e and N_n are the electron and neutron densities in matter respectively. The $\pm$ is positive for neutrinos and negative for antineutrinos. The NC contributions from protons and electrons cancel out in regular matter.

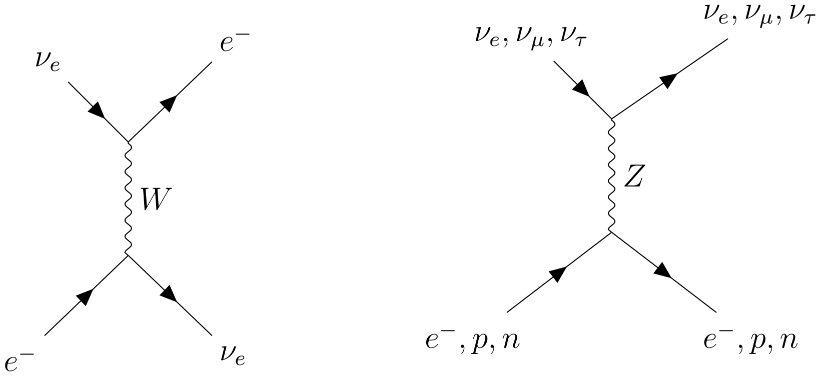

Figure 5.3: Tree level Feynman diagrams that lead to neutrino forward coherent elastic scattering processes in non-exotic matter, showing charged current (CC, left) and neutral current (NC, right) processes [4].

Now, notice that the interaction operator acts on the flavour eigenstates, while the free spacetime translation operator acts on the mass eigenstates, creating a complicated differential equation. Multiplying the left of (5.5) by $\langle\nu_\beta|$, the evolution equation for the transition amplitude $\psi_{\alpha\beta}(x) = \langle\nu_\beta|\nu_\alpha(x)\rangle$ can be computed [4]

$$i\frac{d}{dx}\mathbf{\Psi}_\alpha = H_F\mathbf{\Psi}_\alpha, \tag{5.8}$$

where

$$\mathbf{\Psi}_\alpha = \begin{pmatrix}\psi_{\alpha e}(x)\\ \psi_{\alpha\mu}(x)\\ \psi_{\alpha\tau}(x)\end{pmatrix}, \quad H_F = \frac{1}{2E}\left(U\mathbb{M}^2U^\dagger + \mathbb{A}\right), \tag{5.9}$$

$$\mathbb{M}^2 = \begin{pmatrix}0 & 0 & 0\\ 0 & \Delta m^2_{21} & 0\\ 0 & 0 & \Delta m^2_{31}\end{pmatrix}, \quad \mathbb{A} = \begin{pmatrix}A_{CC} & 0 & 0\\ 0 & 0 & 0\\ 0 & 0 & 0\end{pmatrix}, \quad A_{CC} = \pm 2EV_{CC} = \pm 2\sqrt{2}EG_FN_e. \tag{5.10}$$

Notice that the NC potential was removed since it affects all flavours equally, and is thus an arbitrary re-phasing of neutrino states. Furthermore, it is still only the mass splittings that come into play, not the mass values themselves.

5.3.2 Two-Flavour Approximation

Limiting oneself to electron neutrinos mixing with only one other flavour is useful to gain a qualitative idea the effects this matter potential has. A clear derivation and discussion is presented in Ref. [4], from which a couple main results are laid out here. The neutrino oscillation behaviour with mixing anlge θ and mass splitting Δm^2 depends on the so-called *adiabaticity parameter*

$$\gamma = \frac{\left(\Delta m_M^2\right)^2}{2E\sin\left(2\theta_M\right)\left|dA_{CC}/dx\right|}, \tag{5.11}$$

which is a measure of magnitude of electron density variation compared to the effective mass splitting Δm_M^2, where

$$\begin{aligned} \Delta m_M^2 &= \sqrt{\left(\Delta m^2\cos\left(2\theta\right) - A_{CC}\right)^2 + \left(\Delta m^2\sin\left(2\theta\right)\right)^2}, \\ \tan\left(2\theta_M\right) &= \frac{\tan\left(2\theta\right)}{1 - \frac{A_{CC}}{\Delta m^2\cos(2\theta)}}. \end{aligned} \tag{5.12}$$

If $\gamma \gg 1$, the neutrino evolution is said to be adiabatic. In this case, the electron (anti)neutrino survival probability is simply governed by the standard two-flavour formula

$$P_{\nu_e\to\nu_e}(L,E) = 1 - \sin^2\left(2\theta_M\right)\sin\left(\frac{\Delta m_M^2 L}{4E}\right), \tag{5.13}$$

with the effective mixing angle and mass splitting defined above. From equations (5.12) it is clear that the effective mixing can become maximal $\theta_M = \pi/4$ at a particular resonant electron density obeying $A_{CC} = \Delta m^2\cos\left(2\theta\right)$, which can lead to total flavour conversion if the particular density region is large enough. This is known as the Mikheyev-Wolfenstein-Smirnov (MSW) effect. Such total adiabatic flavour conversion can occur inside the sun (discussed in section 5.4.2), in which case the two-flavour long baseline approximation can be used. Notice that whether this resonance is possible, and its energy dependence, depend on whether θ is above or below $\pi/4$, on the sign of Δm^2, and on whether these are neutrinos or antineutrinos.

If the adiabaticity assumption does not hold, some other approximations are still possible, in particular regimes, as discussed again in [4]. Otherwise, numerically solving the two or three flavour evolution equation is required.

5.3.3 Three-Flavours

There are many ways to approach the full three-flavour problem, which have already been written about extensively. A sample is shown here for context. First, the mass hierarchy can be exploited to separate out the contributions of each mass difference, as described above. The rest of this section is lifted almost verbatim from the author's paper on the topic [10].

If the background matter density varies too quickly, or one mass difference does not totally dominate, one must keep track of all the transition amplitudes. Refs. [68] and [69] use diagonalisation to compare with the vacuum case and determine the effective mass difference and mixing angles in matter. This is the standard approach, including numerical techniques: numerically diagonalise at each iteration of the evolution equation [70, 71]. Meanwhile, Ref. [72] uses Lagrange's formula to determine the evolution operator

$$U(L) = \sum_{n} \left[\prod_{m \neq n} \frac{H_F - \mathcal{E}_m \mathbf{1}}{\mathcal{E}_n - \mathcal{E}_m} \right] \exp\left(-i \frac{\mathcal{E}_n^2 L}{2E} \right), \tag{5.14}$$

where $\mathcal{E}_n$ are the three eigenvalues of H_F, with rather involved expressions provided for constant matter density. The expression is written here explicitly due to its similarity with what will be shown later on. For its part, Ref. [73] uses the Cayley-Hamilton theorem to decompose the evolution operator into a linear combination of second order polynomials of mixing matrices, with analytic expressions for these matrices as well as their coefficients. Recently, Ref. [74] used the eigenvector-eigenvalue identity to derive relatively simple formulae for the effective mixing angles and CP violating phase, along with perturbative approximations of these and mass differences, which can be used in the vacuum expressions as normal. A summary of many of these exact and approximate techniques can be found in Ref. [75], along with very useful accuracy and speed comparisons in the context of long baseline $\overset{(-)}{\nu}_e$ appearance experiments. Ref. [76] provides an elegant generalisation to $(3+N)$ neutrino flavours, and a generic matter potential that can include non-standard interactions (NSI). The effects of the sterile neutrinos and NSI are studied both together and independently. These last two papers also make their codes available via public GitHub repositories, referenced therein.

Lastly, two recent papers [77, 78] compute a general form for the evolution operator (assuming constant matter density) in terms of a Gell-Mann basis and structure constants, using methods from Ref. [73]. The first paper then derives perturbative expansions for particular electron density profiles in the Earth, while the second formulates a general method to compute the oscillation probability of any general time independent Hamiltonian for two or three active neutrino flavours. The initial approach of these will be used as the starting point in the development of a novel algorithm described in section 5.5, also taken from the author's paper. Before that however, the following sections are not from the paper.

5.3.4 Application to SNO+, the Electron Density

Even though the SNO+ detector acts as a long baseline reactor neutrino detector, the full three flavour oscillation formula is used for more accuracy in its oscillation analysis. However, reactor neutrinos necessarily travel through some amount of the Earth's interior before reaching the detector. It is therefore desirable to either improve this three-flavour calculation to include matter effects, or at the least quantify their impact.

This impact depends directly on the traversed electron density, which itself depends on the depths reached by the antineutrinos. To estimate this depth, a spherical Earth is assumed with radius R_E, the SNO+ detector at a depth δ below the surface, and a nuclear reactor at a distance P away from the SNO+ detector's surface location, along the Earth's surface. For such a setup, one can show that the maximum depth reached by a straight line between the rector and the detector is given by

$$d = R_E \left[1 - \frac{(R_E - \delta)\sin\left(\frac{P}{R_E}\right)}{\sqrt{\delta^2 + 4R_E\,(R_E - \delta)\sin^2\left(\frac{P}{2R_E}\right)}} \right]. \tag{5.15}$$

Now over 60% of the reactor antineutrino flux at SNO+ comes from 350 km away or less, and the detector is roughly 2 km underground, for which the maximum depth reached is 3.5 km. This is well within the Earth's crust, which is at least 40km thick around that area [79]. In fact, the furthest a nuclear reactor can be with its neutrinos never going below 40 km in depth is roughly 1410 km. The area encompassed by this is shown in figure 5.4a, including a great deal of the North American reactors, which amount to over 95% of the incoming reactor antineutrino flux at SNO+.

Looking at the crust compositions in [80], its average density is $\rho \approx 2.7$ g/cm^3, and the average ratio of protons to nucleons is $\langle N/A \rangle \approx 0.5$. Using an average nucleon mass of $m_A \approx 1.66 \times 10^{-24}$ g, the electron density can be related to these quantities via

$$N_e = \langle N/A \rangle \frac{\rho}{m_A}, \tag{5.16}$$

so that finally

$$\frac{A_{CC}}{E} \approx 2.05 \times 10^{-7} \left(\frac{\mathrm{eV}^2}{\mathrm{MeV}} \right). \tag{5.17}$$

Variations in this density are rather small [79], especially increases. So for a range of energies between 1 MeV and 10 MeV, and the largest mass splitting around 2.5×10^{-3} eV^2, $A_{CC} \ll \Delta m_{21}, \Delta m_{31}$ is clearly subdominant. The biggest variation might be a drop of up to 2/3 if the neutrinos traverse a significant amount of water. The closest reactor complex, Bruce, is located on the edge of lake Huron, which stands between it and SNO+. However, it turns out to at worst only briefly skim the shallow ($\approx$ 60 m) edge of the lake at the start of the beam, while later the deepest part of the lake is merely 150m [81], while the beam is around its deepest point there. The depth map of lake Huron is shown in figure 5.4b. No other important antineutrino sources appear to cross any significant amount of water either.

Combining the sub-dominance of A_{CC} and the fairly constant crustal density, through which over 95% of incoming antineutrinos pass, a constant lithospheric density is a very good approximation for the SNO+ reactor antineutrino analysis.

5.4 Oscillation Parameter Measurements

The main objective of the present analysis being the measurement of θ_{12} and especially Δm^2_{21}, a brief history of their measurements are written here, followed by a summary of the current best fit values. θ_{13} and Δm^2_{31} are also reviewed to a lesser degree, as they do enter the survival probability calculation, albeit in a less impactful way with regards to this analysis. The last two oscillation parameters (θ_{23} and δ_{CP}) are not discussed here.

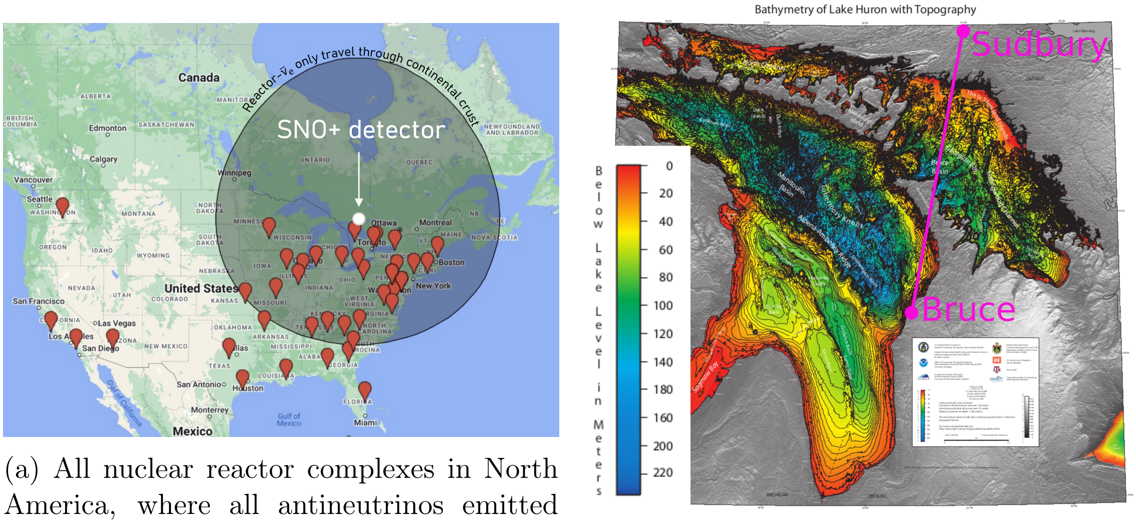

(a) All nuclear reactor complexes in North America, where all antineutrinos emitted from reactors inside the black circle only travel through the Earth's crust before reaching the detector. Uses Google Maps.

(b) Depth Map of Lake Huron [81], with approximate Bruce-Detector Neutrino Path Overlaid.

Figure 5.4: Maps of surrounding nuclear reactors and topology around the SNO+ detector.

5.4.1 Short baseline Experiments

Initial short baseline reactor neutrino experiments in the 1990s, such as Palo Verde [82] in Arizona, and the CHOOZ experiment [83] northern France found no evidence of neutrino oscillation. Their results were compatible with $\theta_{13} = 0$, and thus no $\overline{\nu}_e$ disappearance, with over 90% confidence in some regimes.

Three next generation experiments were designed to finally confirm whether or not θ_{13} is null: the upgraded Double Chooz experiment [84], the RENO experiment [85] in South Korea, and the Daya Bay experiment [86] in China. These were all Gadolinium-doped liquid scintillator experiments, to increase the neutron capture efficiency so critical to tagging IBDs. They also each consisted of multiple detectors, with at the very least one near detector and one far detector roughly 1 km away from the reactor source, so that their results are independent of reactor flux modelling. This distance is optimal to measure short baseline oscillation, as is evidenced by figure 5.2a. They indeed produced extremely precise measurements of θ_{13} through the 2010s, as well as measurements of the atmospheric mass splitting. This was generally written as a measurement of $\left|\Delta m^2_{32}\right|$, due to the previously mentioned sign insensitivity and the equivalence with $\left|\Delta m^2_{31}\right|$. The final results of these three are compatible, and combined in the 2022 PDG table's global fit values of

$$s^2_{13} = (2.20 \pm 0.07) \times 10^{-2}, \tag{5.18}$$

and

$$\Delta m^2_{32} = \begin{cases} (-2.536 \pm 0.034) \times 10^{-3}\text{eV}^2 & \text{(IO)} \\ (2.453 \pm 0.033) \times 10^{-3}\text{eV}^2 & \text{(NO)} \end{cases}, \tag{5.19}$$

for the inverted and normal ordered cases. In general the normal ordered value is used throughout this text unless otherwise stated, though obviously the choice is negligible to this long baseline analysis. These values are thus now precise at the percent, or several percent, level.

The three latter experiments also produced much more accurate measurements of reactor antineutrino fluxes and spectra, informing either directly or indirectly the spectra and models used in chapter 4.

5.4.2 Long baseline Experiments

Solar Neutrinos – Very Long Baseline

Solar neutrino experiments were the first players on the field of neutrino oscillation, discovering the so-called solar neutrino problem [4, 87] via the Homestake experiment [88], confirming it with the Kamiokande [89], GALLEX/GNO [90], SAGE [91] and Super-Kamiokande [12] experiments, and finally solving it in favour of neutrino oscillation with the SNO experiment [24] around the year 2010. These experiments measure the neutrino flux produced by various processes of the nuclear fusion cycle within the sun, as shown in figure 5.5. These are all emitted as electron neutrinos. The solar neutrino review from Ref. [92] provides a good overview of the history and underlying physics of this field, in much more detail than is presented here. In brief, Super-Kamiokande and SNO are/were water Cherenkov detectors, allowing them access to directionality information from elastic scattering (ES) events. SNO used almost exactly the same hardware as its successor SNO+, while Super-K is a 33 kiloton detector, around 1 km underground. This allowed Super-K to confirm the neutrino deficit, and that they indeed came from the Sun, with the first real-time detection of solar neutrinos. SNO was later designed to use heavy water, so that neutrinos might undergo CC and NC reactions with deuterium d too:

- ES: $\nu_x + e^- \to \nu_x + e^-$,
- NC: $\nu_x + d \to \nu_x + n + p$,
- CC: $\nu_e + d \to e^- + p + p$,

where ν_x represents any neutrino flavour. Now while a deficit was detected in ES events due to a deficit of electron neutrinos, which have a 6 times larger ES cross-section compared to other flavours [23] (access to the CC and NC diagrams for ES shown in figure 5.3, rather than just the NC), all neutrino flavours can undergo some form of ES. The CC interaction with deuterium allowed a pure electron neutrino signal to be detected, while its NC interaction provided a completely flavour agnostic neutrino signal. In this way, the electron neutrinos were shown to oscillate to other flavours without doubt. Until just this year, this provided the only model-independent solar neutrino flux. Initial results from the PandaX-4T [93] and XENONnT [94] dark matter search experiments used neutrino coherent elastic scattering on xenon nuclei to reject the background-only hypothesis in favour of a ^{8}B flux by almost 3σ each.

Solar neutrino experiments are very sensitive to θ_{12} due to the oscillation baseline, as one would expect. From the averaging out of oscillation at such long distances seen for example in figure 4.6 and equation (5.4), one might expect them to be relatively insensitive to Δm^2_{21}. However, they turn out to be somewhat sensitive to it due to matter effects, both in the Sun and in the Earth, as the neutrinos travel across high density regions. In the Earth, the different flux detected between day and night, due to neutrinos traveling very different distances in the Earth, is mainly

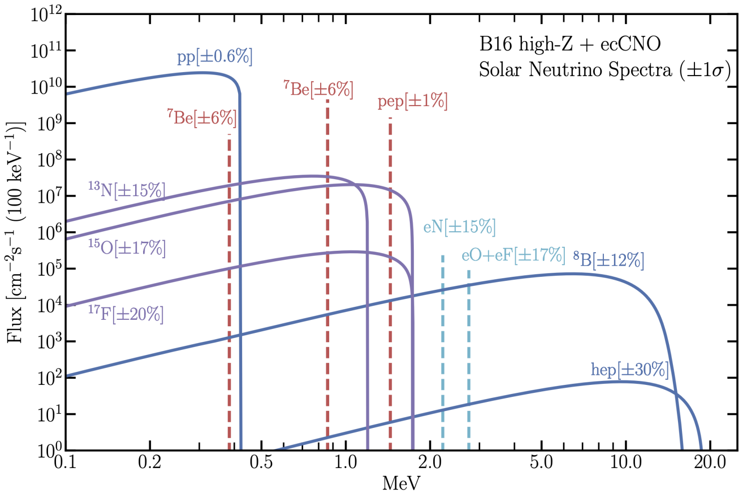

Figure 5.5: Emitted solar neutrino spectra from processes in the nuclear fusion cycles [92].

affected by Δm^2_{21}, from which an absolute value of it was measured by Super-K and SNO [95].

In addition, at the densities inside the Sun, neutrinos above around 5 MeV can pass through the MSW resonance discussed in section 5.3.2, undergoing large scale adiabatic flavour transition. This effect is energy dependent, but Super-K and SNO had energy thresholds on the MeV scale – only giving them access to the ^{8}B flux. Borexino [96] achieved a far lower energy threshold and backgrounds level, allowing measurement of the pp, ^{7}Be and pep fluxes too, and thus the MSW effect at different energy scales [97]. The fractions of these measured fluxes to those predicted by the standard solar model produce survival probabilities of electron neutrinos at the different energies of these fluxes, as shown in figure 5.6.

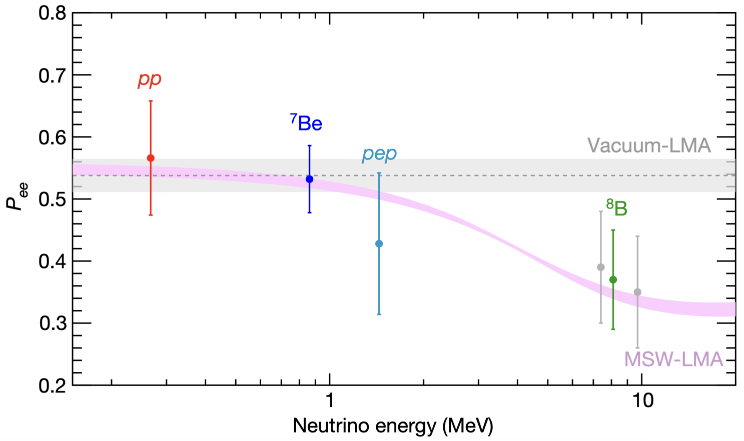

Figure 5.6: Electron neutrino survival probability measured at different energies, from different solar neutrino flux measurements by Borexino: the pp, ^{7}Be, pep and ^{8}B fluxes. The ^{8}B measurement is displayed in low and high energy ranges (grey) and a combined measurement (green). The grey horizontal line shows the expected survivial probability in a vacuum (Vacuum - Large Mixing Angle), while the pink curve is the prediction including the MSW effect. They both display $\pm 1\sigma$ bands. Figure taken from Ref. [97].

This shows that at lower energies, the survival probability is consistent with the average vacuum value computed previously from equation (5.4), while at higher energies a decrease in the survival probability is observed. This can be understood by first observing that for such long baseline oscillation, only Δm^2_{21} and especially θ_{12} matter. Second, recall how the effective mixing angle of such two-flavour oscillation depends on the intervening matter density and the neutrino energy via equation (5.12). As described in that section, the survival probability would decrease if $\Delta m^2_{21} > 0$ and $\theta_{12} < \pi/4$, or if $\Delta m^2_{21} < 0$ and $\theta_{12} > \pi/4$, in the way sketched out by the pink curve in figure 5.6. These measurements are thus sensitive to the sign of the Δm^2_{21} mass splitting, determining it to be positive.

An overall solar neutrino oscillation fit was recently updated by the Super-K collaboration, based mainly on theirs and SNO's data, yielding [95]

$$\begin{aligned}\Delta m^2_{21} &= 6.10^{+0.95}_{-0.81} \times 10^{-5}\text{eV}^2,\\ s^2_{12} &= 0.306 \pm 0.013.\end{aligned} \tag{5.20}$$

Reactor Neutrinos – Medium Baseline

Until the SNO+ results being released in the last few months [3] (and concurrent publication to appear around the time of this thesis' submission), the only long baseline neutrino oscillation measurement from nuclear reactors was by the KamLAND experiment, at the Kamioka observatory in Japan. The KamLAND experiment is similar to SNO+, being a ~1 kiloton liquid scintillator detector, surrounded by PMTs, 1 km underground. A diagram of KamLAND's overall design is presented in figure 5.7, showing the 13m diameter balloon filled with the liquid scintillator, itself surrounded by mineral oil to act as shielding. The PMTs surround it all at about a 9m radius from the centre, achieving roughly 34% photocathode coverage. Being 1km underground, the detector receives a muon rate of roughly 0.3Hz, and it is surrounded by roughly 50 nuclear reactor cores in Japan, with a weighted average distance of around 180 km [98].

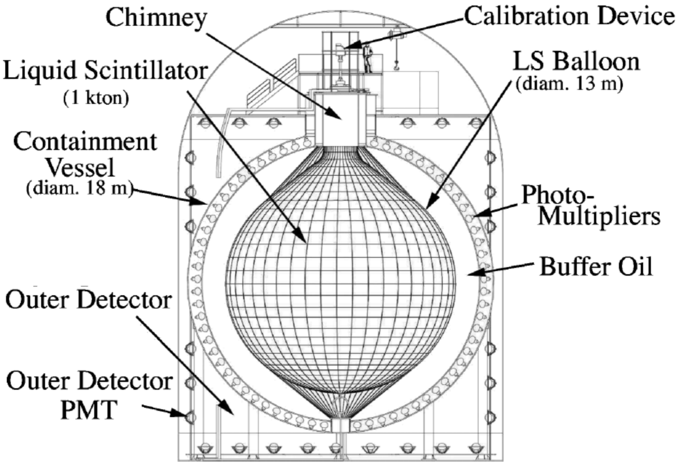

Figure 5.7: KamLAND experimental design [98].

After roughly 7 years of data taking, KamLAND's results are as follows [99]

$$\begin{aligned}\Delta m^2_{21} &= (7.53 \pm 0.18) \times 10^{-5}\text{eV}^2,\\ \tan^2\theta_{12} &= 0.436^{+0.029}_{-0.025},\end{aligned} \tag{5.21}$$

as well as the confirmation $s^2_{13} = 0.023 \pm 0.002$. At present, this is the most accurate measurement of Δm^2_{21} by any method, and disagrees slightly with the solar results, as will be shown in the next section. The result for θ_{12} on the other hand agrees with solar results very well, improving its measurement.

With the larger detector volume and more close-by reactors, the reactor IBD rate at KamLAND is around 1-2 per day [99], roughly three to four times higher than at SNO+. The PMT coverage of KamLAND is only slightly less than at SNO+, and their energy resolution is similar. One would expect KamLAND to be far more sensitive to Δm^2_{21}, both from its higher flux and its closer location to reactors. However, the distribution of reactors around SNO+ gives it a particular advantage. Most of the flux at KamLAND comes from many reactors over different distances, such as 79% from 26 cores between 138 and 214 km away, and 6.7% from one reactor 88km away, which averages out some of the spectral impact of the oscillation. Conversely, around 40% of the flux at SNO+ comes from a single powerful reactor complex 240km away (Bruce), while around 20% more from around 350km away, with other major sources significantly further away. As such, SNO+ hopes to surpass KamLAND's measurement within the next couple of years, as will be discussed in the final oscillation analysis results of section 8.4.2.

5.4.3 Current best measurements

The most up to date oscillation parameters – used in the analysis herein, unless otherwise stated – are thus summarised in table 5.1, taken from the 2022 PDG booklet [6]. The Δm^2_{21} value is simply taken from KamLAND, as the most accurate and in slight tension with the solar results. Other parameters are obtained from global fits such as NUFIT, though s^2_{12} mainly comes from the combined solar fit, and s^2_{13} and Δm^2_{32} from the short baseline reactor experiments.

Table 5.1: Oscillation parameters used throughout this analysis, unless otherwise stated. Taken from PDG 2022 [6], with the normal ordered value for Δm^2_{32}.

Δm^2_{21}	$(7.53 \pm 0.18) \times 10^{-5}$ eV^2
Δm^2_{32}	$(2.453 \pm 0.033) \times 10^{-3}$ eV^2
s^2_{12}	0.307 ± 0.013
s^2_{13}	$(2.20 \pm 0.07) \times 10^{-2}$

The slight ($\sim 1.5\sigma$) tension between the combined solar fit and the KamLAND result is shown in figure 5.8a. Using the best fit oscillation parameters just discussed, and both of these Δm^2_{21}, the expected reactor IBD spectrum at SNO+ is shown in figure 5.8b for each case. It is clear that SNO+ should be quite sensitive in distinguishing between these two results. The KamLAND result is the one used from now on for any predictions or comparison, as stated previously. In this case, the not oscillated IBD flux from the previous chapter (assuming perfect detection efficiency and 100% reactor outputs) drops from roughly 63 to 35 per year.

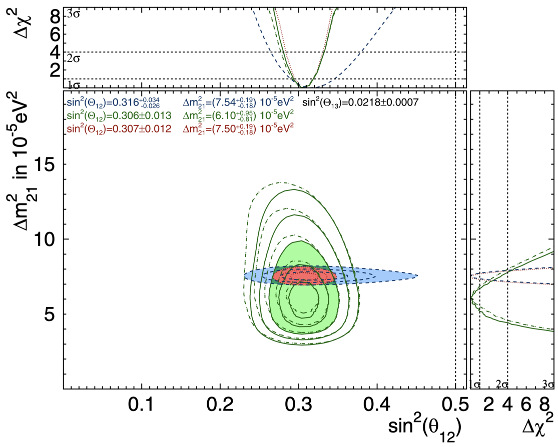

(a) Long baseline oscillation allowed parameter space from global fits. The solid green lines are 1 to 5 σ contours from Super-K's combined global solar analysis, while the nearby dashed-dotted lines use only Super-K and SNO data. The dashed blue lines show the 1 to 3 σ contours of the KamLAND result, while the red dotted lines in the middle are the contours the of combined result. Shaded areas highlight the 3-σ contours in each case [95].

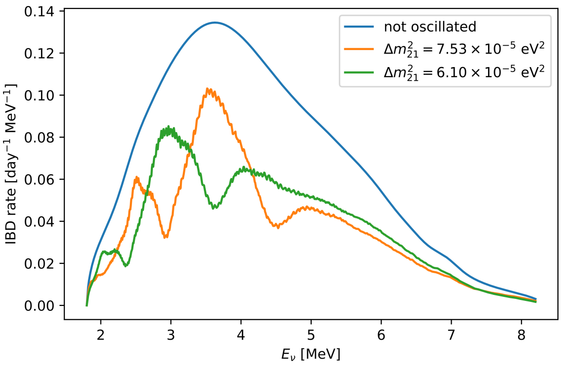

(b) Expected reactor IBD spectrum at SNO+, taken from figure 4.6, with neutrino oscillation applied. Two different values for Δm^2_{21} are used to show the sensitivity to this parameter.

Figure 5.8: The solar versus KamLAND Δm^2_{21} tension's impact on the expected reactor IBD spectrum at SNO+.

5.5 New Matter Effects Algorithm

As per the above discussion, a formula for three flavour neutrino oscillation assuming constant matter density is desired for SNO+. In particular, it must lead to an algorithm that is relatively fast if it is to be used in an oscillation analysis. This is since oscillation parameters must be allowed to float, and the survival probability re-computed a great many times. Such an algorithm was derived and published in Ref. [10] by the author of this thesis, and so most of the paper is copied into this section, almost verbatim due to its direct relevance and use in this analysis.

5.5.1 Solving the Differential Equation

The Evolution Operator

This subsection largely follows Bushra Shafaq and Faisal Akram's method in [77]. First, the traceless effective Hamiltonian H is defined, since the trace acts on all flavours equally and so does not contribute to mixing

$$\begin{aligned} &H \equiv H_F - \frac{1}{3}\mathrm{tr}\left[H_F\right]\mathbf{1}, \\ &\mathrm{tr}\left[H_F\right] = \Delta m_{21}^2 + \Delta m_{31}^2 + A_{CC}. \end{aligned} \tag{5.22}$$

The evolution equation is thus

$$i\frac{d}{dx}\mathbf{\Psi}_\alpha = \frac{1}{2E}H\mathbf{\Psi}_\alpha, \tag{5.23}$$

where contrary to Shafaq and Akram's paper, the $1/2E$ factor is kept separate from H. Assuming constant matter density, this is solved with

$$\mathbf{\Psi}_\alpha(x) = U(x)\mathbf{\Psi}_\alpha(0),\ \ U(x) = \exp\left(-iH\frac{x}{2E}\right), \tag{5.24}$$

where $U(x)$ is the evolution operator. Now, these can be decomposed using the property that the Gell-Mann matrices ($\lambda^i,\ i \in \{1, ..., 8\}$) and the identity matrix form a complete orthogonal basis for 3×3 complex matrices. H is traceless, so it does not need the identity matrix

$$\begin{aligned} &H = h^i\lambda^i,\ h^i = \frac{1}{2}\mathrm{tr}\left[H\lambda^i\right], \\ &U(x) = u_0\mathbf{1} + iu_i\lambda^i,\ u_0 = \frac{1}{3}\mathrm{tr}\left[U(x)\right],\ u_i = \frac{1}{2i}\mathrm{tr}\left[U(x)\lambda^i\right], \end{aligned} \tag{5.25}$$

where from now on repeated dummy indices imply summation. These equations are derived from the Gell-Mann matrix identities $\mathrm{tr}\left[\lambda^i\lambda^j\right] = 2\delta_{ij}$ and $\mathrm{tr}\left[\lambda^i\right] = 0$.

Now, some useful general results in linear algebra will be used: for a matrix A, with eigenvalues $\mathcal{E}[A]_n$,

$$\begin{aligned} &\det(A) = \prod_n \mathcal{E}[A]_n,\ \mathrm{tr}\left[A\right] = \sum_n \mathcal{E}[A]_n, \\ &\text{and if } B = f(A) \text{ and } f \text{ is a holomorphic function,} \\ &\mathcal{E}[B]_n = f(\mathcal{E}[A]_n). \end{aligned} \tag{5.26}$$

Therefore, recalling $U(x) = \exp\left(-iHx\right)$ and defining $\mathcal{E}[H]_n = \mathcal{E}_n$,

$$u_0 = \frac{1}{3}\sum_{n=0}^{2} \exp\left(-i\frac{\mathcal{E}_n x}{2E}\right). \tag{5.27}$$

For u_i, first note that from $H = h^i\lambda^i$,

$$\frac{\partial U(x)}{\partial h^i} = -\frac{it}{2E}\lambda^i U(x), \tag{5.28}$$

and so using the previous identities one can show

$$u_i = \frac{-i}{2}\sum_{n=0}^{2}\frac{\partial\mathcal{E}_n}{\partial h^i}\exp\left(-i\frac{\mathcal{E}_n x}{2E}\right). \tag{5.29}$$

All that is needed now are expressions for the eigenvalues $\mathcal{E}_n$ of H. The parametric equation of a 3×3 matrix A with eigenvalues λ is

$$\begin{aligned}\det\left(A-\lambda\mathbf{1}\right) =& -\lambda^3 + \mathrm{tr}(A)\lambda^2 - \frac{1}{2}\left(\mathrm{tr}(A)^2 - \mathrm{tr}\left(A^2\right)\right)\lambda \\ &+ \det(A) \\ =& 0,\end{aligned} \tag{5.30}$$

so that for the traceless H,

$$\begin{aligned}&\mathcal{E}_n^3 - 3a_1\mathcal{E}_n - 2a_0 = 0,\\ &a_1 = \frac{1}{6}\mathrm{tr}[H^2] = \frac{1}{3}h^i h^i,\\ &a_0 = \frac{1}{2}\det\left(H\right) = \frac{1}{3}d^{ijk}h^i h^j h^k,\end{aligned} \tag{5.31}$$

where d^{ijk} are the symmetric structure constants of the Gell-Mann matrices

$$\begin{aligned}\{\lambda^i,\lambda^j\} &= \frac{4}{3}\delta_{ij}\mathbf{1} + 2d^{ijk}\lambda^k,\\ d^{ijk} &= \frac{1}{4}\mathrm{tr}\left(\lambda^i\{\lambda^j,\lambda^k\}\right).\end{aligned} \tag{5.32}$$

The last relation between the determinant and structure constants in (eqn 5.31) can be derived by first multiplying the structure constant definition (eqn 5.32) (second equation) by $h^i h^j h^k$ (and summing over these indices as normal), to find

$$\mathrm{tr}\left(H^3\right) = 2d^{ijk}h^i h^j h^k. \tag{5.33}$$

Then from the definition of the determinant of a 3×3 matrix

$$\det(H) = \frac{1}{3!}h^i h^j h^k \epsilon_{a_1a_2a_3}\epsilon_{b_1b_2b_3}\lambda^i_{a_1b_1}\lambda^j_{a_2b_2}\lambda^k_{a_3b_3}, \tag{5.34}$$

the Levi-Civita identity $\epsilon^{a_1a_2a_3}\epsilon_{b_1b_2b_3} = 3!\delta^{[a_1}_{b_1}\delta^{a_2}_{b_2}\delta^{a_3]}_{b_3}$ (the index position is irrelevent here), and recalling that the Gell-Mann matrices are traceless ($\lambda^i_{aa} = 0$), one can find

$$\det(H) = \frac{1}{3}\mathrm{tr}\left(H^3\right), \tag{5.35}$$

and thus

$$\det(H) = \frac{2}{3} d^{ijk} h^i h^j h^k. \tag{5.36}$$

Meanwhile, taking the derivative of the parametric equation (eqn 5.31) w.r.t h^i gives the needed expression

$$\frac{\partial \mathcal{E}_n}{\partial h^i} = \frac{2}{3} \frac{h^i \mathcal{E}_n + d^{ijk} h^j h^k}{\mathcal{E}_n^2 - a_1}, \tag{5.37}$$

while different solutions to the parametric equation are used here compared to the original paper

$$\mathcal{E}_n = 2\sqrt{a_1} \cos \left[\frac{1}{3} \cos^{-1} \left(\frac{a_0}{a_1^{3/2}} \right) - \frac{2\pi n}{3} \right], \quad n \in \{0, 1, 2\}. \tag{5.38}$$

These are the solutions to a depressed cubic equation for real solutions, which must be real since H is Hermitian (one can also check that $a_0, a_1 \in \mathbb{R}$, and $\frac{a_0^2}{4} + \frac{a_1^3}{27} < 0$, which imply the solutions are real). The evolution operator is then

$$\mathcal{U}(x) = \frac{1}{3} \sum_{n=0}^{2} \left(1 + \frac{\mathcal{E}_n H + Y}{\mathcal{E}_n^2 - a_1} \right) \exp \left(-i \frac{\mathcal{E}_n x}{2E} \right), \tag{5.39}$$

with

$$Y \equiv d^{ijk} h^i h^j \lambda^k = H^2 - 2a_1 \mathbf{1}, \tag{5.40}$$

which can be shown from the first equation of (eqn 5.32) and multiplied by $h^i h^j$ (summing over indices). This last relation (eqn 5.40) was not in Bushra Shafaq and Faisal Akram's paper [77].

Finally, assuming a (anti)neutrino is produced in a pure flavour state $\psi_{\alpha\beta}(0) = \delta_{\alpha\beta}$ and recalling $P_{\nu_\alpha \to \nu_\beta}(x) = \left| \psi_{\alpha\beta}(x) \right|^2$, one therefore has the transition probability

$$\begin{aligned} P_{\nu_\alpha \to \nu_\beta}(L, E) =& \sum_{n,m} (X_n)_{\beta\alpha} (X_m)^*_{\beta\alpha} \exp \left[-i \frac{(\mathcal{E}_n - \mathcal{E}_m) L}{2E} \right], \\ X_n =& \frac{1}{3} \left(\mathbf{1} + \frac{\mathcal{E}_n H + Y}{\mathcal{E}_n^2 - a_1} \right), \end{aligned} \tag{5.41}$$

where $x = L$ is the propagation length, as usual. This equation is of course of the same form as the vacuum case, but writing out the effective mass differences and mixing angles is saved for a later section.

Details and Simplifications in Vacuum

The rest of this derivation departs from [77], and is entirely original work. Notice that variable quantities such as L and E only appear in the last expression (eqn 5.41), except for where E and N_e enter into A_{CC} at the beginning. Because of the structure of H in terms of A_{CC}, it will turn out that most calculations can be performed with vacuum settings ($A_{CC} = 0$), and small modifications added later to take into account matter effects (see the next section). Therefore, here we take a look at the details assuming a vacuum first, where all the associated quantities will be marked with a tilde for clarity $H_F = \tilde{H}_F + \mathbb{A}$.

Here $\tilde{H}_F$ is simply $\tilde{H}_F = U\mathbb{M}U^\dagger$, and from the cyclic nature of the trace $\mathrm{tr}[\tilde{H}_F] = \mathrm{tr}[\mathbb{M}]$, so that

$$\begin{aligned} \tilde{H} &= U\mathbb{M}U^\dagger - \frac{1}{3}\mathrm{tr}[\mathbb{M}]\mathbf{1}, \\ \mathrm{tr}[\mathbb{M}] &= \Delta m^2_{21} + \Delta m^2_{31}, \end{aligned} \tag{5.42}$$

and thus the components are explicitly given by

$$\tilde{H}_{\alpha\beta} = \sum_{f=2,3} \Delta m^2_{f1} \left(U_{\alpha f} U^*_{\beta f} - \frac{1}{3}\delta_{\alpha\beta} \right). \tag{5.43}$$

Now, $\tilde{a}_1$ and $\tilde{a}_0$ can be computed from $\tilde{h}^i$ and d^{ijk}, but it is easier to use the definitions $\tilde{a}_1 = \frac{1}{6}\mathrm{tr}\left[\tilde{H}^2\right]$ and $\tilde{a}_0 = \frac{1}{2}\mathrm{det}\left(\tilde{H}\right)$. For $\tilde{a}_1$ it is straightforward to show, using (eqn 5.42)

$$\tilde{a}_1 = \frac{1}{9}\left[(\Delta m^2_{21})^2 + (\Delta m^2_{31})^2 - \Delta m^2_{21}\Delta m^2_{31} \right], \tag{5.44}$$

while for $\tilde{a}_0$, the formula (eqn 5.30) for $\det(A - \lambda\mathbf{1})$ can be reused, with $A = U\mathbb{M}U^\dagger$ and $\lambda = \frac{1}{3}\mathrm{tr}[\mathbb{M}]$, so that

$$\begin{aligned} \tilde{a}_0 =& \frac{1}{27}\left[(\Delta m^2_{21})^3 + (\Delta m^2_{31})^3 \right] \\ &- \frac{1}{18}\left[(\Delta m^2_{21})^2 \Delta m^2_{31} + \Delta m^2_{21} (\Delta m^2_{31})^2 \right], \end{aligned} \tag{5.45}$$

where use was made of $\det\left(U\mathbb{M}U^\dagger\right) = \det(U)\det(\mathbb{M})\det(U^\dagger)$ and $\det(\mathbb{M}) = 0$. Lastly, one can show that

$$\tilde{Y}_{\alpha\beta} = \frac{1}{3}\sum_{f=1}^{3} \left(\Delta m^2_{f1}\right)^2 \left(U_{\alpha f} U^*_{\beta f} - \frac{1}{3}\delta_{\alpha\beta} \right), \tag{5.46}$$

where $\left(\Delta m^2_{11}\right)^2 \equiv 2\Delta m^2_{21}\Delta m^2_{31}$ is defined for compactness. So for a vacuum, these quantities can all be substituted in to compute $\mathcal{E}_n$ (eqn 5.38), X_n and $P_{\nu_\alpha \to \nu_\beta}(L)$ (eqn 5.41) directly. Notice also that the eigenvalues $\mathcal{E}_n$ only depend on the mass differences here, as one would expect.

Adding Matter Effects

The values calculated above must be corrected for matter effects. From $H_F = \tilde{H}_F + \mathbb{A}$, the traceless matrix H can be related to the vacuum one $\tilde{H}$ simply according to A_{CC}

$$H = \tilde{H} + \frac{1}{3}A_{CC}D, \quad D = \begin{pmatrix} 2 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{pmatrix}. \tag{5.47}$$

Corrections to Y are also easier to see in matrix notation

$$Y = \tilde{Y} + \frac{1}{3}A_{CC}T + \frac{1}{9}A^2_{CC}D, \tag{5.48}$$

$$T = \begin{pmatrix} 2\tilde{H}_{ee} & \tilde{H}_{e\mu} & \tilde{H}_{e\tau} \\ \tilde{H}^*_{e\mu} & 2\tilde{H}_{\tau\tau} & -2\tilde{H}_{\mu\tau} \\ \tilde{H}^*_{e\tau} & -2\tilde{H}^*_{\mu\tau} & 2\tilde{H}_{\mu\mu} \end{pmatrix}. \tag{5.49}$$

Since only the diagonal components of $\tilde{H}$ change, from $a_1 = -\frac{1}{2}\mathrm{tr}(H^2)$ one can find

$$a_1 = \tilde{a}_1 + \frac{1}{3}\tilde{H}_{ee}A_{CC} + \frac{1}{9}A^2_{CC}, \tag{5.50}$$

and using the determinant definition of a_0, it is modified by

$$\begin{aligned} a_0 =&\tilde{a}_0 + \frac{1}{6}A_{CC}\left(\tilde{H}^2_{ee} + 2\tilde{H}_{\mu\mu}\tilde{H}_{\tau\tau} - 2|\tilde{H}_{\mu\tau}|^2 \right.\\ &\left. + |\tilde{H}_{e\mu}|^2 + |\tilde{H}_{e\tau}|^2\right) + \frac{1}{6}A^2_{CC}\tilde{H}_{ee} + \frac{1}{27}A^3_{CC}, \end{aligned} \tag{5.51}$$

which one can find is simply

$$a_0 = \tilde{a}_0 + \frac{1}{2}\tilde{Y}_{ee}A_{CC} + \frac{1}{6}\tilde{H}_{ee}A^2_{CC} + \frac{1}{27}A^3_{CC}. \tag{5.52}$$

Notice that since the diagonal components of $\tilde{H}$ are real, so are those of $\tilde{Y}$, and therefore a_0 and a_1 and, by extension $\mathcal{E}_n$, are always real. X_n consequently always has real diagonal components, as expected.

5.5.2 Example Algorithm: Electron (Anti)Neutrino Survival Probability

What turns out to be the simplest example of how this can all be used in an algorithm is shown here: the electron (anti)neutrino survival probability. Other example algorithms are also shown in the original paper [10], but are not relevant here. The calculation is composed of two steps: the first performed once to compute some constant values, and the second using these values for each particular (anti)neutrino energy and/or electron density.

First one should compute the following four constant quantities, written here in terms of the mass differences and mixing angles of the standard PMNS matrix parametrisation:

$$\tilde{H}_{ee} = \Delta m^2_{21}\left(s^2_{12}c^2_{13} - \frac{1}{3}\right) + \Delta m^2_{31}\left(s^2_{13} - \frac{1}{3}\right), \tag{5.53}$$

$$\begin{aligned} \tilde{Y}_{ee} = \frac{1}{3}\Bigg[& \left(\Delta m^2_{21}\right)^2\left(s^2_{12}c^2_{13} - \frac{1}{3}\right) \\ & + \left(\Delta m^2_{31}\right)^2\left(s^2_{13} - \frac{1}{3}\right) \\ & + 2\Delta m^2_{21}\Delta m^2_{31}\left(c^2_{12}c^2_{13} - \frac{1}{3}\right)\Bigg], \end{aligned} \tag{5.54}$$

$$\begin{aligned} \tilde{a}_0 =& \frac{1}{27}\left[(\Delta m^2_{21})^3 + (\Delta m^2_{31})^3\right] \\ & - \frac{1}{18}\left[(\Delta m^2_{21})^2\Delta m^2_{31} + \Delta m^2_{21}(\Delta m^2_{31})^2\right], \end{aligned} \tag{5.55}$$

and

$$\tilde{a}_1 = \frac{1}{9}\left[(\Delta m^2_{21})^2 + (\Delta m^2_{31})^2 - \Delta m^2_{21}\Delta m^2_{31}\right]. \tag{5.56}$$

Then the next part is performed for a given electron (anti)neutrino energy E and matter electron density N_e. First A_{CC} is computed:

$$A_{CC} = \pm 2\sqrt{2}G_F E N_e, \tag{5.57}$$

then the constants corrected for this:

$$H_{ee} = \tilde{H}_{ee} + \frac{2}{3}A_{CC}, \tag{5.58}$$

$$a_0 = \tilde{a}_0 + \frac{1}{2}\tilde{Y}_{ee}A_{CC} + \frac{1}{6}\tilde{H}_{ee}A^2_{CC} + \frac{1}{27}A^3_{CC}, \tag{5.59}$$

$$a_1 = \tilde{a}_1 + \frac{1}{3}\tilde{H}_{ee}A_{CC} + \frac{1}{9}A^2_{CC}, \tag{5.60}$$

$$Y_{ee} = \tilde{Y}_{ee} + \frac{2}{3}\tilde{H}_{ee}A_{CC} + \frac{2}{9}A^2_{CC}. \tag{5.61}$$

These can then be substituted into

$$\mathcal{E}_n = 2\sqrt{a_1}\cos\left[\frac{1}{3}\cos^{-1}\left(\frac{a_0}{a_1^{3/2}}\right) - \frac{2\pi n}{3}\right], \quad n \in \{0, 1, 2\}. \tag{5.62}$$

$$(X_n)_{ee} = \frac{1}{3}\left(1 + \frac{\mathcal{E}_n H_{ee} + Y_{ee}}{\mathcal{E}_n^2 - a_1}\right), \tag{5.63}$$

so that finally

$$P_{\overset{(-)}{\nu_e}\to\overset{(-)}{\nu_e}} = 1 - 4\sum_{n>m}(X_n)_{ee}(X_m)_{ee}\sin^2\left((\mathcal{E}_n - \mathcal{E}_m)\frac{L}{4E}\right). \tag{5.64}$$

The fact that the components $(X_n)_{ee}$ are real was used to derive the more compact formula (eqn 5.64). Notice that neither θ_{23}, nor δ_{13} enter this calculation, no matter the electron density. The (anti)electron neutrino survival probability is thus independent of these for constant densities.

5.5.3 Results

Speed Comparison for Example Algorithm

In order to get a sense of the speed of this algorithm, calculations of various neutrino oscillation probabilities in constant matter density were performed by the above algorithms, written in C++. The same was then done with the widely used General Long Baseline Experiment Simulator (GLoBES) [70] – a flexible software package also written in C++, allowing the simulation of the propagation of neutrinos from a user-defined source to a user-defined detector via any matter density profile. It additionally includes an analysis framework, to provide results and sensitivities of the simulated experiment to different oscillation channels, accounting for a potentially wide array of user-defined systematic uncertainties and constraints. Critically, it includes a library of neutrino oscillation functions, such as *glbConstantDensityProbability()*, which computes the transition or survival probability between any two neutrino flavours for constant matter density, neutrino energy and baseline. It does this by diagonalising the Hamiltonian with various numerical or analytic methods, depending on which is fastest [73, 71]. It is a fast and reliable method, but all matrix elements must be recomputed for each change in neutrino energy E, baseline L and matter density ρ, while this paper's method must only re-perform part of the calculations. Therefore, while the GLoBES function was not fully optimised for speed – being part of a more general framework – it is a good benchmark test, due to its wide usage in the neutrino community. The results and processing times of this function were thus compared to the algorithm herein.

Method First, an initialisation step is performed, where the pre-computed constants above are calculated and GLoBES is initialised. This step was not timed since it need only be performed once and so will not scale with the number of calculations. However, note that the GLoBES initialisation takes longer since the package includes many more functionalities than just constant matter neutrino oscillations.

Second, for a given flavour transition, and a range of 100 neutrino energies, 100 baselines and 100 matter densities, the oscillation probabilities were computed using three different functions: the GLoBES function, a general flavour version of this algorithm (Section II), and then the version of this algorithm tailored to the specific flavour transition (Section III). The results and total computation times (CPU time as measured by the C++ standard library *std::clock()* function) of these three were recorded. This process was repeated 50 times to obtain a measure of statistical uncertainty.

The ranges of neutrino energies E, baselines L and matter densities ρ were evenly spaced values in some range

$$\begin{aligned} &E_{\text{min}} \leq E \leq E_{\text{max}}, \\ &L_{\text{min}} \leq L \leq L_{\text{max}}, \\ &\rho_{\text{min}} \leq \rho \leq \rho_{\text{max}}, \end{aligned} \tag{5.65}$$

where the minima were always fixed ($E_{\text{min}} = 0.5$ MeV, $L_{\text{min}} = 0.01$ km, $\rho_{\text{min}} = 0$ g/cm^3). Therefore, for a given set of maxima (E_{max}, L_{max}, ρ_{max}), each function was called 50 million times ($100 \times 100 \times 100 \times 50$).

Results As alluded to, the whole process was performed for various maximum values, to discern any E, L or ρ dependence on the results. Ten different values for each were used, according to

$$\begin{aligned} &E_{\text{max}} \in \{x \ : \ E_{\text{min}} \leq x \leq 1000 \text{ MeV}\}, \\ &L_{\text{max}} \in \{x \ : \ L_{\text{min}} \leq x \leq 1000 \text{ km}\}, \\ &\rho_{\text{max}} \in \{x \ : \ \rho_{\text{min}} \leq x \leq 100 \text{ g/cm}^3\}, \end{aligned} \tag{5.66}$$

so that the whole method above was carried out one thousand times.

The computed probabilities were always exactly the same (having copied any unit conversion factors from the GLoBES code), so are not shown here. However, the computation times are presented in figure 5.9 for three example oscillations. Dependence on the three E_{max}, L_{max} and ρ_{max} parameters is shown separately. The plot showing E_{max} dependence averages over all L_{max} and ρ_{max} dependence, and likewise for the other two cases. Statistical uncertainty in these averages was propagated throughout, and added quadratically to a systematic error of $\sigma_{\text{sys}} = 0.01$ s, being the resolution limit of the CPU time measuring function.

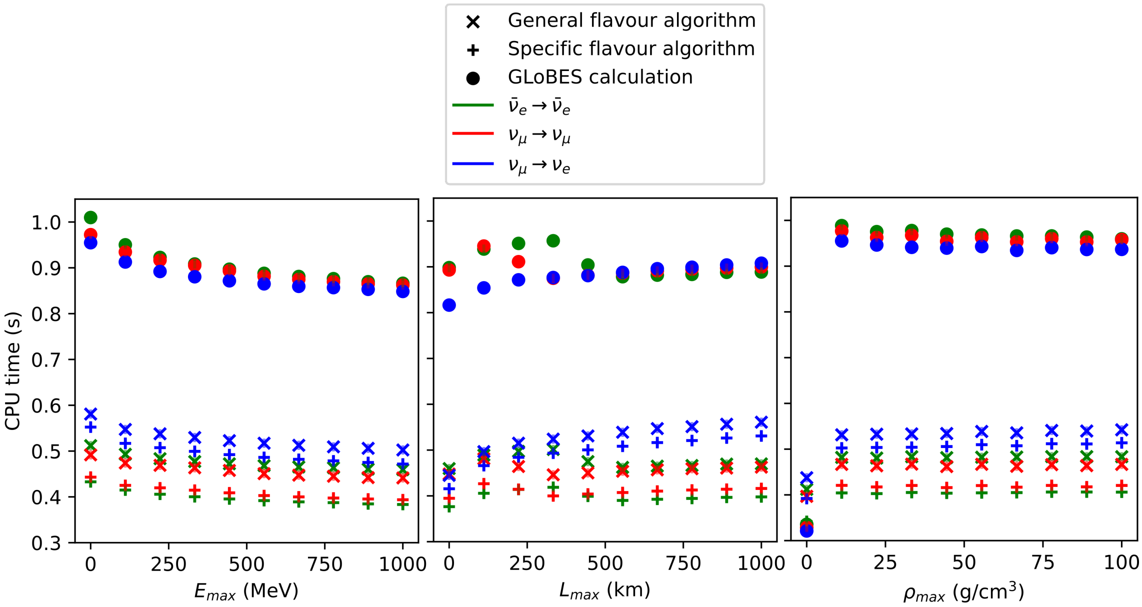

Figure 5.9: Comparison of the CPU time taken by new neutrino oscillation algorithms against the same GLoBES calculation, showing its their dependence on energy (E), baseline (L) and matter density (ρ). Error bars are generally too small to be visible.

For zero matter density, the GLoBES calculation appears slightly faster, while every other configuration shows these to be almost or around twice as fast as GLoBES. To be specific, if T_{GLoBES} is the CPU time taken by the GLoBES function averaged over all the data above, and likewise T_{general} and T_{specific} are the CPU times taken by the general and specific flavour algorithms from this paper respectively,

$$\frac{T_{\text{GLoBES}}}{T_{\text{general}}} = 1.82 \pm 0.04, \quad \frac{T_{\text{GLoBES}}}{T_{\text{specific}}} = 1.98 \pm 0.05. \tag{5.67}$$

Most of the uncertainty comes from the variability in GLoBES' computation time.

Additionally, following discussion with Peter Denton, the $\overset{(-)}{\nu}_e$ appearance algorithm was compared to other calculations described in [75]. Unlike GLoBES, these algorithms are optimised for speed, and in general also tailored to the constant matter density case, meaning they provide a more suitable speed comparison. Use was made of Peter Denton's code on github, where a fork was made (`https://github.com/Jamicus96/Nu-Pert-Compare`) to add this paper's algorithm in two separate cases:

- The first step, or initialisation of quantities in vacuum, is performed separately, before any speed comparison (branch "compare_JP_precomp").

- The initialisation is included in the speed comparison (branch "compare_JP").

It was found that the former was marginally faster than the fastest exact calculation in Peter Denton's code package "ZS" (as of February 2024) – on the order of 6% faster – while the latter slightly slower – around 17% slower. It was estimated that accounting for the initialisation time, this paper's algorithm becomes faster than the "ZS" algorithm after 3 loops (3 probability calculations for different neutrino energies). The differences here are small, and one must bear in mind that some approximate solutions described in [75], and included in the code, are significantly faster.

Effective Parameters

From the correspondence between the PMNS matrix and X_n matrices found earlier (eqn 5.41), as well as that between the eigenvalue differences and the mass differences, one can find the effective parameters (including matter effects)

$$\begin{aligned} \widehat{\Delta m^2_{kj}} &= \mathcal{E}_m - \mathcal{E}_n, \\ \widehat{U}_{\alpha k}\widehat{U}^*_{\beta k} &= (X_n)_{\alpha\beta}\,, \end{aligned} \tag{5.68}$$

for some relationship between (k, j) and (m, n) indices. To deduce this relationship, note that (5.68) must hold for the vacuum case, so the relationship need only be shown for that simplified case. Now, the eigenvalues of the traceless matrix $\tilde{H}$ (5.22) in the vacuum case are clearly

$$\begin{aligned} \lambda_1 &= -\frac{1}{3}\left(\Delta m^2_{21} + \Delta m^2_{31}\right), \\ \lambda_2 &= \frac{2}{3}\Delta m^2_{21} - \frac{1}{3}\Delta m^2_{31}, \\ \lambda_3 &= -\frac{1}{3}\Delta m^2_{21} + \frac{2}{3}\Delta m^2_{31}, \end{aligned} \tag{5.69}$$

so the vacuum eigenvalues $\tilde{\mathcal{E}}_n$ must be assigned to these in some order. Next, looking at the definition of $\mathcal{E}_n$ (5.38), the $\frac{1}{3}\cos^{-1}(...)$ term is always between 0 and $\frac{\pi}{3}$. This means that

$$\mathcal{E}_0 \;>\; \mathcal{E}_1 \;>\; \mathcal{E}_2, \tag{5.70}$$

and $\mathcal{E}_0 \;>\; 0$ always hold. Therefore, the ordering of λ_1, λ_2 and λ_3 determines their relationship. This ordering depends on the mass ordering itself: for Normal Ordering (NO) $\lambda_3 > 0 > \lambda_2 > \lambda_1$, and for Inverted Ordering (IO) $\lambda_2 > \lambda_1 > 0 > \lambda_3$. Therefore, the relationship between the indices in (5.68) is shown in table 5.2. For example in normal ordering, $\widehat{\Delta m^2_{21}} = \mathcal{E}_1 - \mathcal{E}_2$ and $|\widehat{U}_{e3}|^2 = (X_0)_{ee}$. Thus, the effective

(k, j)	NO (m, n)	IO (m, n)
3	0	2
2	1	0
1	2	1

Table 5.2: Index correspondence for effective parameters from equation 5.68, for Normal Ordering (NO) and Inverted Ordering (IO).

mixing angles can be evaluated in terms of these too, such as in normal ordering for example:

$$\begin{aligned}
\widehat{s_{13}}^2 &= (X_0)_{ee}\,, \\
\widehat{s_{12}}^2 &= \frac{(X_1)_{ee}}{\widehat{c_{13}}^2}, \\
\widehat{s_{23}}^2 &= \frac{(X_0)_{\mu\mu}}{\widehat{c_{13}}^2}, \\
\cos\widehat{\delta_{13}} &= \frac{\widehat{c_{12}}^2\widehat{c_{23}}^2 + \widehat{s_{12}}^2\widehat{s_{23}}^2\widehat{s_{13}}^2 - (X_1)_{\mu\mu}}{2\widehat{c_{12}}\widehat{c_{23}}\widehat{s_{12}}\widehat{s_{23}}\widehat{s_{13}}}.
\end{aligned} \tag{5.71}$$

Notice that any dependence on energy or electron density in these parameters comes only from factors of A_{CC}. Their values can thus be plotted on a simple graph against $A_{CC} \propto EN_e$, without having to vary E and N_e independently. Additionally, negative values of A_{CC} can be used to plot the behaviour of antineutrinos, since flipping the sign of A_{CC} is effectively the only difference. See figure 5.10a for the fractional scaling of these parameters in the MeV neutrino energy scale (for lithospheric electron density). To see large scale absolute changes, such as mass differences changing ordering, one must go to the GeV scale, as shown in figure 5.10b. As one might expect, “regime changes" (sudden changes in evolution of effective parameters) appear when A_{CC} reaches the same scale as the mass differences. Notice also that $\widehat{\Delta m^2_{21}}$ and $\widehat{s_{12}}^2$ are the most strongly affected by matter in the MeV scale.

Furthermore, one can use these effective oscillation parameters in short and long baseline approximate formulae, just as one would with the vacuum constants. For example the long baseline approximation

$$P = 1 - \frac{1}{2}\sin^2(2\theta_{13}) - \sin^2(2\theta_{12})\, c_{13}^4 \sin^2\left(\frac{L\Delta m^2_{21}}{4E}\right), \tag{5.72}$$

becomes, for Normal Ordering,

$$P = 1 - 4\,(X_1)_{ee}\,(X_2)_{ee} - 2\,(X_0)_{ee}\sin^2\left(\frac{L\left(\mathcal{E}_1 - \mathcal{E}_2\right)}{4E}\right), \tag{5.73}$$

which follows the full matter effect oscillation formula more closely.

5.6 SNO+ IBD Spectrum

Departing from the paper [10] for good, it is clear from figure 5.8b that the expected IBD spectrum is very sensitive to Δm^2_{21}, while the small scale oscillations due to Δm^2_{31} are washed out by all the different reactor sources, and the energy resolution of the detector. Recall the short baseline wavelength over L/E is $4\pi/\Delta m^2_{31} \approx 1\ \mathrm{km}\cdot\mathrm{MeV}^{-1}$, from which the wavelength over just energy at a particular baseline is given by $\lambda(E) = (4\pi/\Delta m^2_{31})\,(E^2/L)$. Thus the energy wavelength for antineutrinos from Bruce (240 km away) around the IBD's 4 MeV peak is $\lambda(E) \approx 0.07$ MeV, as shown in figure 5.11a.

Figure 5.11a compares the vacuum and matter effects survival probabilities for antineutrinos coming from Bruce, for both full three-flavour and two-flavour long baseline approximations. One can see that the matter effect long baseline approximation tracks its three-flavour counterpart well. Therefore, to compare the impact of matter effects on the expected IBD spectrum at SNO+, the long baseline formulae (5.2) and (5.73) are used for clarity in figure 5.11b. Meanwhile, the integrated flux from all four combinations are shown in table 5.3, where it is clear that matter effects induce a larger change than the use or not of the long baseline approximation, though still only by around 0.6%. Meanwhile, spectral changes are induced by modifications to the effective $\widehat{\Delta m^2_{21}}$ parameter on the order of 0.5 to 1%, as shown in figure 5.10a.

Table 5.3: Comparison of expected number of reactor IBD events per year at SNO+ using different survival probability calculations. Uses un-oscillated numbers as described in figure 4.6, and 2-flavours implies the use a long baseline approximation. Arbitrary accuracy is used in the numbers simply to showcase the minor difference in predictions. For reference, using $\langle P_{\overline{\nu}_e \to \overline{\nu}_e} \rangle$ yields 32.39.

	3-flavour	2-flavour
Vacuum	32.128	32.130
Matter	32.328	32.326

5.7 Summary

The SNO+ detector acts as a very sensitive long baseline oscillation detector for reactor antineutrinos, thanks to the layout and power of nearby nuclear reactors. A novel three-flavour matter-effect algorithm, and associated two-flavour (long baseline) approximation were derived. The effect of matter is a 0.6% increase in the IBD flux, with sub-percent level changes to the spectrum. Using the long baseline approximation changes the flux negligibly, and only eliminates small scale oscillation that is of the order of the detector's energy resolution or less. The full three-flavour matter effect oscillation formula is used regardless, since it is only very marginally more computationally demanding than the vacuum case. Meanwhile this algorithm ensures the highest accuracy into the future as increasingly large data samples are collected, and other systematic uncertainties are potentially reduced.

Following the publication of this algorithm, a separate comparison of some of the fastest such calculations – including the present one – was carried out in Ref. [100], and one of its plots is displayed in figure 5.12. The "aggressive" compiler option `-Ofast` referenced therein activates a few other compiler options, such as `-O3` and `-ffast-math`. The former enables slew of compiler optimisations, which reduce the execution time at the expense of larger compilation time and reduced debugging facility. This option is rather safe. On the other hand, `-ffast-math` allows for some mathematical shortcuts, such as reduced error handling (ignoring the possibility of "nan" values), assuming commutativity in operations, and so on. These further reduce execution time, and while they are mostly safe for simple and thoroughly tested use cases, they can lead to incorrect results for programs that depend on an exact implementation of IEEE or ISO specifications for maths functions (technical standard for floating-point arithmetic). Documentation on these compiler flags can be found in Ref. [101], while a helpful discussion and warning on `-ffast-math` is presented in Ref. [102].

Much of the speed improvement for the author's algorithm ("Page") visible in figure 5.12 could likely be achieved through judicious operation reordering and other code optimisations, before compilation. These considerations aside, the algorithm achieves speeds almost on par with Peter Denton's newly optimised "NuFast" package, themselves nearly as fast as the vacuum calculation. Meanwhile, GLoBES is significantly slower, though as previously stated, it is a more general package not fully optimised for this calculation. Therefore the algorithm presented in this work does not aim to supplant GLoBES, which fulfills a much larger role, but is intended to be easily incorporated into existing oscillation analysis frameworks that lack a fast oscillation probability calculation in constant matter density. In this way it can provide the speed needed for these analysis computations, without needing to install an entire code package such as GLoBES, which is comparatively cumbersome and slow to initialise. The NuFast package presented Ref. [100] appears to have a similar purpose.

For the purposes of error propagation however, the vacuum long baseline approximation is evidently accurate enough to be of use, and provides an easier and more intuitive understanding. Looking again at equation (5.2), the only oscillation parameters that enter, besides the two parameters of interest (θ_{12}, Δm^2_{21}), is θ_{13}. The uncertainty in s^2_{13} is $\sigma_{13} = 0.0007$, as stated previously 5.1, which propagates to an uncertainty in the survival probability as

$$\sigma_P\left(L, E\right) = \frac{2}{|c_{13}|}\left(1 - P(L, E)\right)\sigma_{13}, \tag{5.74}$$

where $P(L, E) = P_{\overline{\nu}_e \to \overline{\nu}_e}(L, E)$. Now, the survival probability is always between 0 and 1, and $c_{13} \approx 1$, so the uncertainty from this is always less than or roughly equal to 0.07%.

The only other quantity from this formula that can contribute uncertainty is the baseline L (E will be included in the energy scaling systematic uncertainties). First, due to isotropic emission, the flux is proportional to $1/L^2$, and so its contribution is a fractional uncertainty of $2\sigma_L/L$. Conservatively assuming a baseline of order $L \approx \mathcal{O}\left(100\text{km}\right)$, and an uncertainty of around $\sigma_L \approx \mathcal{O}\left(100\text{m}\right)$, representing roughly the scale of reactor buildings, this contributes a 0.1% uncertainty to the flux at most. Secondly, this can impact the measured oscillation wavelength via the energy spectral shape, and by extension Δm^2_{21}. Recall the $4\pi/\Delta m^2_{21} \approx 33\ \text{km} \cdot \text{MeV}^{-1}$ wavelength, which for a representative 4 MeV antineutrino is 132 km. The $\mathcal{O}\left(100\text{m}\right)$ uncertainty is thus less than 0.1% of it. This argument also ignores the fact that multiple spectral features (two dips and peaks for example) help reduce the dependence on an absolute baseline prediction.

Lastly, changes and uncertainty in the electron density are considered negligible, since they are subdominant to the already almost negligible constant matter effects. As a result, nothing in this chapter – from neutrino oscillation to the baselines – is considered to contribute uncertainty to the present analysis.

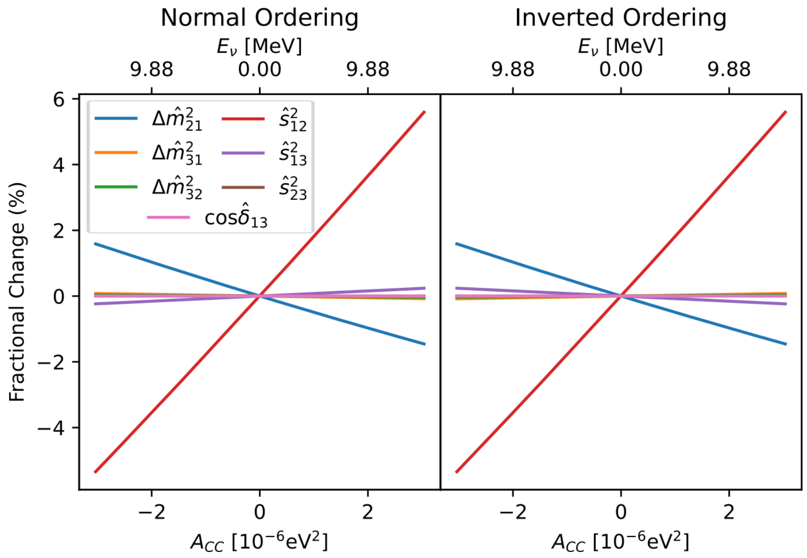

(a) MeV scale (for $\rho = 2.7\text{g/cm}^3$ and $\langle N/A \rangle = 0.5$).

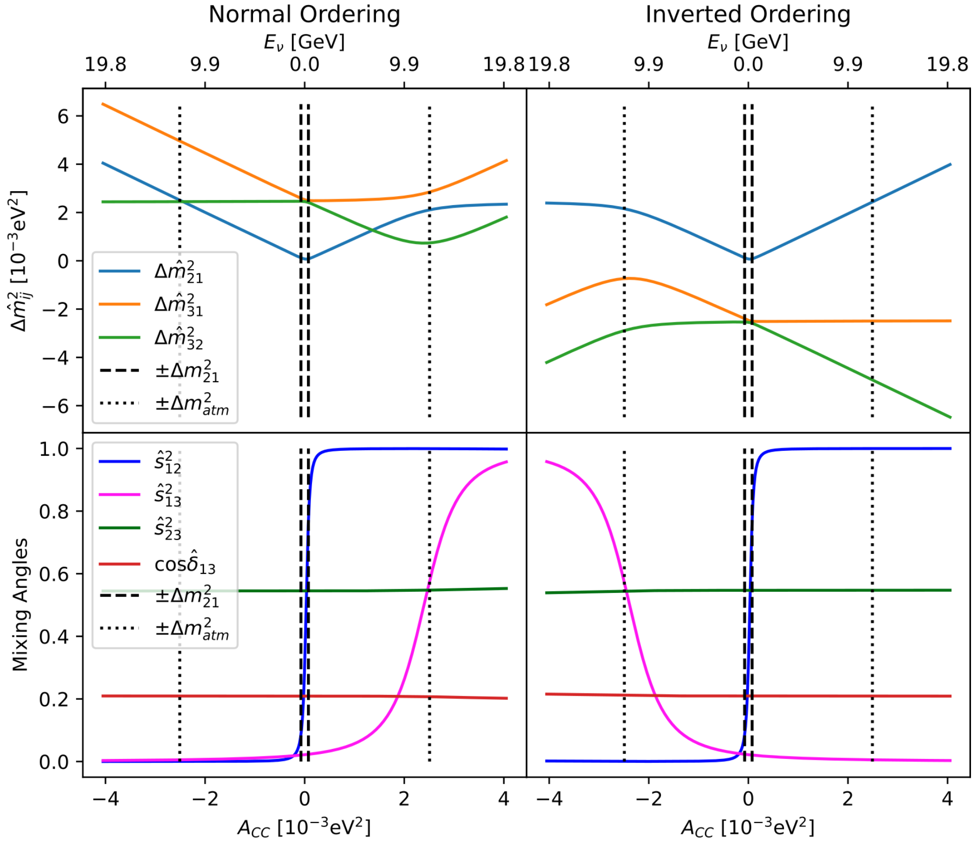

(b) GeV scale (for $\rho = 2.7\text{g/cm}^3$ and $\langle N/A \rangle = 0.5$).

Figure 5.10: Change of the effective neutrino oscillation constants with A_{CC}. Positive values of A_{CC} represent neutrinos, while negative values represent antineutrinos. Vertical lines denote where A_{CC} is equal to the vacuum squared mass-differences, and specifically $\Delta m^2_{\text{atm}} \equiv \frac{1}{2}\left(\Delta m^2_{32} + \Delta m^2_{31}\right)$ is the average atmospheric oscillation mass-difference.

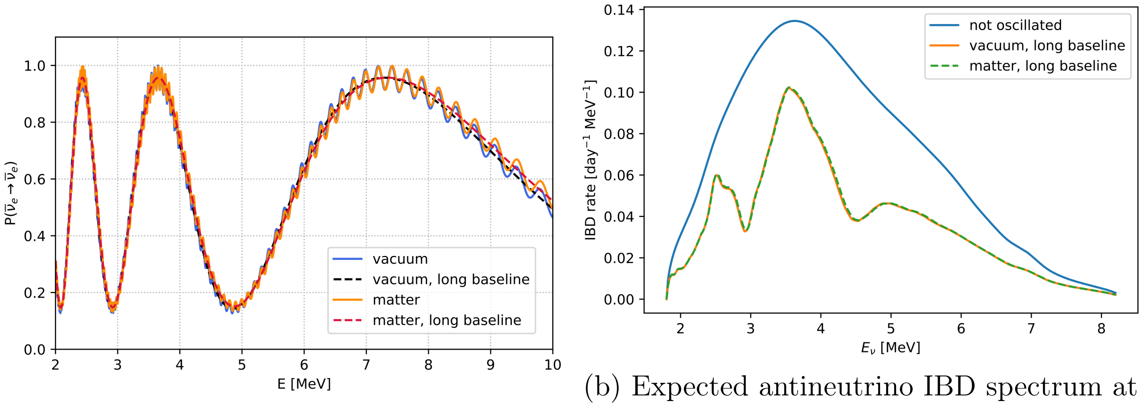

(a) Energy dependence of the survival probability of antineutrinos from the Bruce reactor complex.

(b) Expected antineutrino IBD spectrum at SNO+, using long baseline approximations and the un-oscillated spectrum from figure 4.6.

Figure 5.11: Comparing neutrino Oscillation at SNO+ with and without matter effects.

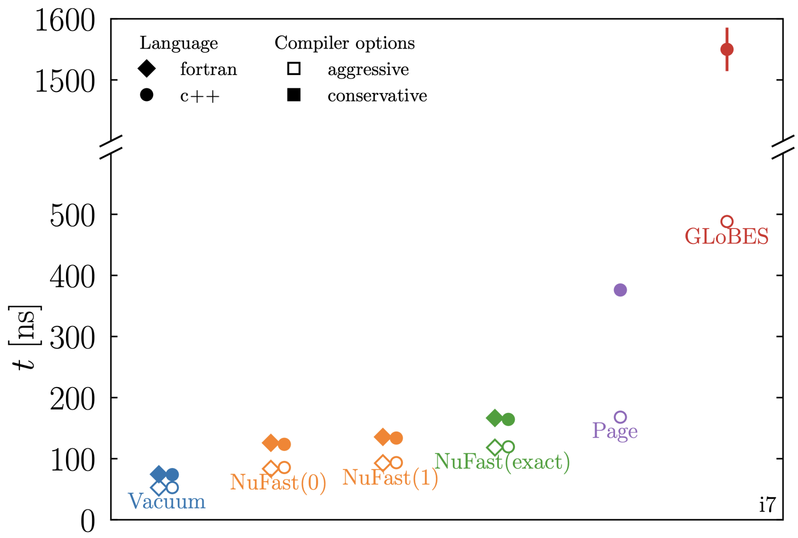

Figure 5.12: The computation time for the calculation of nine neutrino oscillation probabilities, using one laptop core, for various algorithms, including the one presented in this thesis (Page). Each one is performed with aggressive (`-Ofast` and `-ffast-math`) and conservative (`-O0`) compilation flags. Plot taken from Ref. [100].

Chapter 6

Backgrounds

> "I don't know half of you half as well as I should like; and I like less than half of you half as well as you deserve."
>
> Bilbo Baggins – *The Lord of the Rings*, J.R.R. Tolkien

6.1 Introduction

One of the biggest determining factors of the construction of SNO+ is the need to reduce background levels in the region of interest (ROI, roughly 2.42 to 2.56 MeV) enough to detect the rare potential $0\nu\beta\beta$ decays. Such efforts are arguably the most critical factors in low energy or rare event searches such as this. Even if events are detected, backgrounds place an irreducible upper limit on the accuracy with which such measurements can be made – a limit that higher sample sizes alone cannot overcome. Many of these backgrounds can likewise affect the reactor antineutrino analysis, though in different ways. Firstly, the expected IBDs produce prompt events in a broader range of energies compared to the $0\nu\beta\beta$ ROI. Critically, the prompt-delay structure allows event pair tagging, which greatly reduces almost all backgrounds that do not mimic this structure to an almost negligible level, as will be discussed shortly. This is in a sense the "super power" of IBDs, and thus reactor antineutrinos, which facilitated them as the first neutrinos ever measured [103].

Relevant backgrounds are therefore listed in what follows, divided into uncorrelated and correlated backgrounds. This latter category includes the few relevant backgrounds resistant to coincidence tagging: neutron capture events, which include IBDs from non reactor sources, and (α, n) events. Recall that the IBDs fall broadly in the 0.5 to 10 MeV range for prompt events, with delayed events distributed around 2.2 MeV. Only background events falling into this range need be considered. Regardless of their category, background reducing methods are detailed for each one, with their final detection and tagging efficiencies computed and summarised. With this information in hand, predictions for the number and distributions of events expected to be detected are provided, and finally compared to events tagged from real SNO+ data. Systematic uncertainties are considered and quantified throughout.

6.2 Uncorrelated Backgrounds

Uncorrelated backgrounds do not mimic the IBD's prompt-delayed structure, but two such unrelated events can be tagged as a candidate pair "by accident", simply by virtue of their inherently high rates. Nevertheless, this accidental rate will turn out to be very small after all the cuts described below, and as such all uncorrelated backgrounds are grouped together as "accidentals". First, however, an introduction to the radioactive decay chains responsible for a large portion of these uncorrelated backgrounds, and the more troublesome (α, n) background treated later, is presented for context.

6.2.1 Internal Radioactivity

There are two decay chains that SNO+ must contend with in its search for $0\nu\beta\beta$ decays: the uranium and thorium chains, shown in figure 6.1. These can produce events within the ROI, and so decay rates at various points are monitored closely.

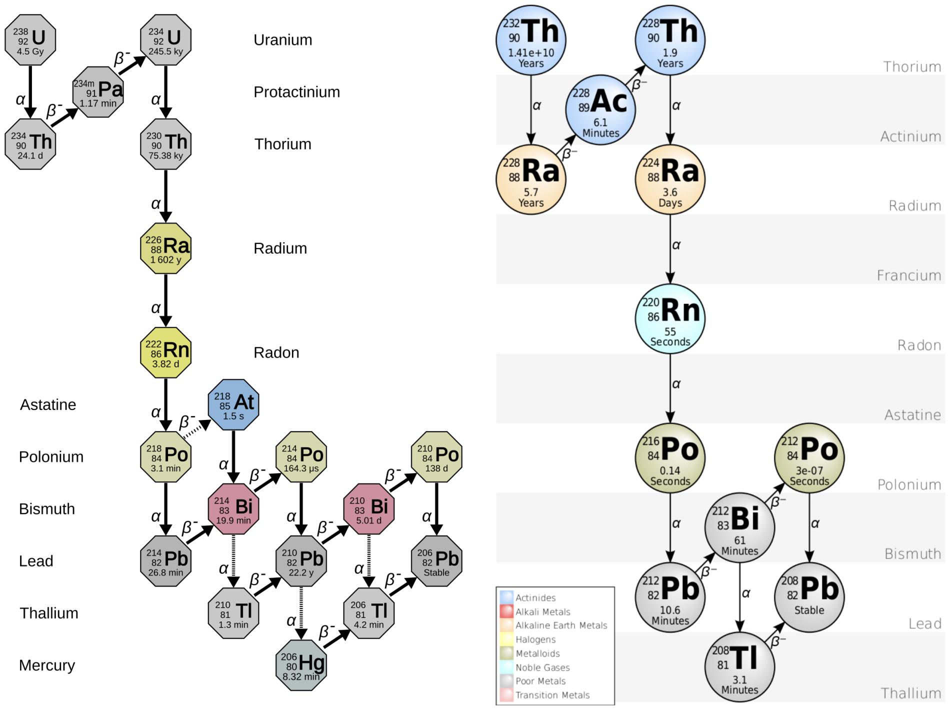

(a) ^{238}U decay chain [104]. (b) ^{232}Th decay chain [105].

Figure 6.1: Primary concerning decay chains at SNO+.

For the purposes of the antineutrino analysis, the main concern is ^{210}Po decays in the uranium chain, since these are the primary source of α emissions in the detector, which can lead to (α, n) events. All other sources of α particles are comparatively negligible. The reason is that the mine air is relatively rich in ^{222}Rn, with a concentration of 123 Bq/m^3. During refurbishment, the AV was exposed to this air, allowing ionic forms of Rn to deposit on the AV surface down to depths of around 100nm [7]. As can be seen in the uranium decay chain, some fraction of this ^{222}Rn can fairly quickly decay to ^{210}Pb through various steps, which has a half-life of 22.3 years. ^{210}Pb is thus still impacted on the AV surface, with a surface

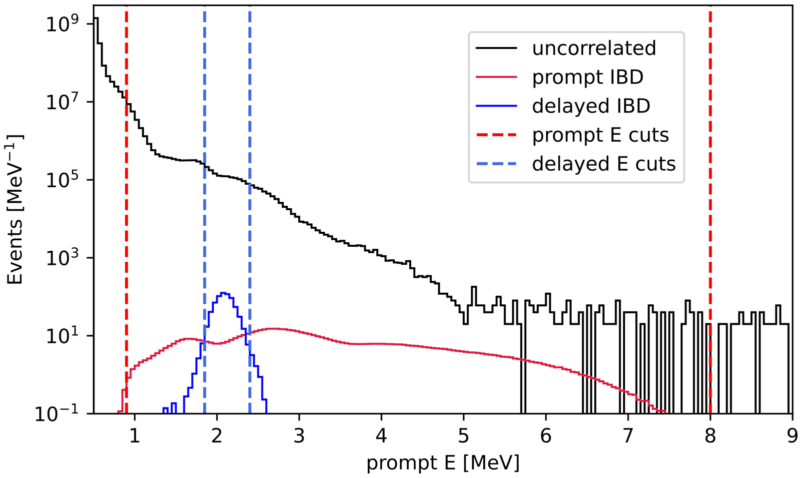

Figure 6.2: Event rates over reconstructed energy. The uncorrelated events are any events measured in the 30/04/2022 and 05/03/2023 gold dataset, excluding muon veto time (see section 6.4), that reconstructed correctly. The IBD events are simulated and oscillated reactor antineutrino induced IBDs, with the expected flux and spectrum in the same period. Prompt and delayed energy cuts are also displayed.

activity measured in 2013 to be $\mathcal{A}_{\text{surface}} = (2.4 \pm 0.8)$ Bq/m^2 [7]. This ^{210}Pb can then produce ^{210}Po via ^{210}Bi, which then emits so-called surface α particles. The ^{210}Pb and daughters can also leach into the liquid scintillator before decaying, with rates on the order of $\lambda = \mathcal{O}\left(10^{-4}\right)$ per day, depending on the scintillator temperature [7], defined according to

$$\frac{d\mathcal{A}_{\text{leached}}(t)}{dt} = \mathcal{A}_{\text{surface}}\lambda e^{-\lambda t}. \tag{6.1}$$

A smaller rate of α particles are thus also produced throughout the liquid scintillator.

The ^{214}BiPo decays described in section 3.7.2 are also products of the uranium decay chain, at far higher rates than IBDs are expected to occur. Thankfully, the ^{214}BiPo's delayed events are almost never in the delayed IBD's energy range. All these radioactive decays are thus grouped together with all other uncorrelated backgrounds.

6.2.2 Background Reduction – Accidental Rate

First, note that only events with a reconstructed radial position of less than 5.7 m are considered here and in everything that follows, due to surface (α, n) events which will be discussed later. Looking at all such events that reconstruct in the broad 0.5 to 9 MeV range in the dataset of interest, an event rate of 8.50 Hz is found. This "uncorrelated" spectrum is shown in figure 6.2, where it is clear that it completely overwhelms the reactor IBD event rate. However, one can reduce this to a small "accidental" background as follows.

First, notice one can cut all events above 8 MeV and below 0.9 MeV with only very minor loss to the prompt IBD rate, while the uncorrelated rate explodes below this lower limit. These prompt cuts reduce the uncorrelated rate to the prompt-like rate

$$r_{\text{prompt}} = 0.124 \text{ Hz}. \tag{6.2}$$

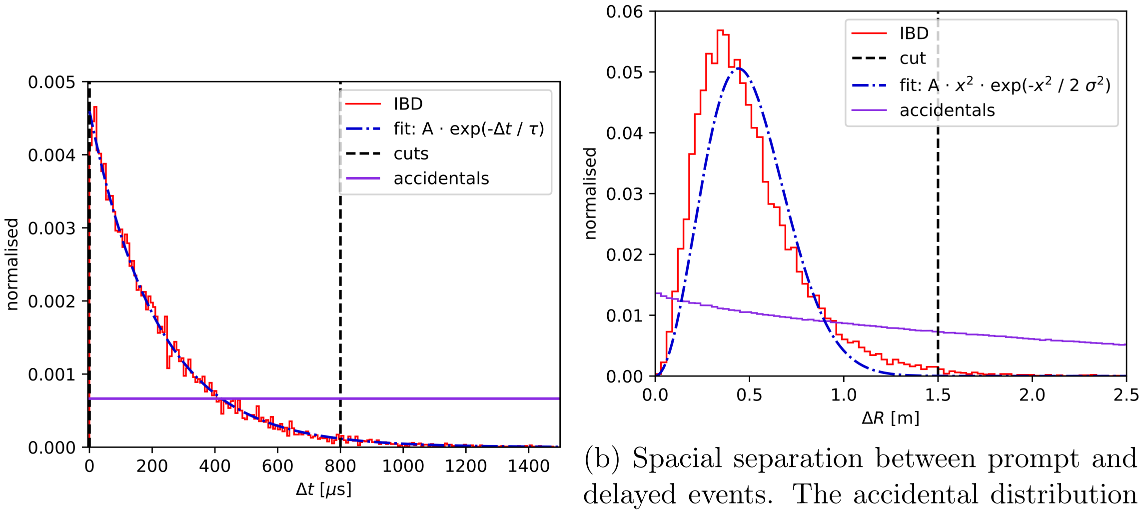

(a) Temporal separation between prompt and delayed events. Constant accidental distribution superimposed, normalised to the range of the plot. An exponential decay was fit to the IBD distribution, yielding a lifetime of $\tau = 217$ μs.

(b) Spacial separation between prompt and delayed events. The accidental distribution was made with a toy MC model simulating events uniformly in a 5.7m spherical volume, and recording distances between these. A random walk displacement PDF with fixed $\sigma = \sqrt{\pi/2}\langle x\rangle$ was fit to the IBD distribution.

Figure 6.3: Reactor antineutrino IBDs compared to accidental coincidences. IBD distributions are from simulated reactor IBD events, and functions are fit to these simply for demonstration proposes.

However, the delayed events provide the lion's share of the discrimination power, being centered around 2.2 MeV. One can thus impose similar cuts around this, at 1.85 and 2.4 MeV, producing an uncorrelated delayed-like event rate of only

$$r_{\text{delayed}} = 6.08 \times 10^{-3} \text{ Hz}, \tag{6.3}$$

while only cutting out a few percent of delayed IBD events.

Now, each delayed IBD event can provide one with its associated prompt event by taking advantage of the time correlation between the two. The time separation between them is governed by an exponential decay with a lifetime of approximately 220 ns, as shown in figure 6.3a. Therefore, one can select only event pairs where the first one passes the above prompt cuts, while the second passes the delayed cuts, and which are separated by less than 800 μs. This is the so-called "coincidence tagging" that has been referenced repeatedly above. Also cutting events separated by less than 0.4 μs, due to the re-triggers mentioned in section 3.4, this reduces the IBD rate by about 2.0% from Δt cuts. Meanwhile, the accidentally coincident rate is given by

$$r_{\text{acc}} = r_{\text{prompt}} \cdot r_{\text{delayed}} \cdot \Delta t_{\text{cut}}, \tag{6.4}$$

where $\Delta t_{\text{cut}} = (800 - 0.4)$ μs. This accidental rate is thus $r_{\text{acc}} = 6.03 \times 10^{-7}$ Hz, finally lower than the expected IBD rate.

One can further reduce this background by noting that the neutron's random walk produces a distinct spacial separation distribution between the IBD prompt and delayed events as shown in figure 6.3b. This distribution is roughly similar to that of a perfect random walk in 3-D, where the mean $\langle \Delta R \rangle$ is related to the number of collisions n and mean free path λ by $\langle \Delta R \rangle = \sqrt{\frac{2n}{\pi}}\lambda$. For the mean $\langle \Delta R \rangle = 0.5$ m in the above figure, and 7 collisions as mentioned in section 6.3.3 below, this gives a λ on the order of 10 cm. Therefore, requiring that a tagged prompt-delayed event pair be separated by less than 1.5 m reduces the IBD flux by less than 1%, while a simple toy MC simulation shows that it reduces randomly paired events in a 5.7 m spherical volume by 91.5%. The final accidental rate is thus

$$r_{\mathrm{acc}} = 5.12 \times 10^{-8}\ \mathrm{Hz}, \tag{6.5}$$

over 7 orders of magnitude lower than the original uncorrelated event rate. This is the strength of IBD coincidence tagging. A more careful consideration of the ΔR cut efficiency can be performed by randomly pairing prompt-like and delayed-like events in real data, thus accounting for the fact that events are not in general uniformly distributed. Such studies reduce the efficiency from 8.5% down to around 5%, though this has not been reviewed in detail by the author of the present analysis, and so the 8.5% efficiency is kept instead, with a 40% uncertainty on the accidentals normalisation. A summary of the cut efficiencies of these and other cuts discussed below will be presented at the end of the chapter.

Lastly note that these cuts were optimised for the antineutrino analysis during the partial-fill phase [106]. They are largely carried over here due to the considerations being very similar. However, one can take advantage of the correlation between the Δt and ΔR distributions, along with prior estimated event rates, to build a likelihood ratio (LR) test between accidental and IBD events. Such a method was employed in [3]; the results of both techniques are similar, and so can be easily compared.

6.3 Correlated Backgrounds

6.3.1 Atmoshperic Backgrounds

Atmospheric Neutrinos

Cosmic rays interacting with particles in the high atmosphere produce a significant flux of neutrinos in the 0.1-10 GeV energy range [4, 107]. Being unimpeded by the rock overburden, they can interact in the SNO+ detector through various means. The $\overline{\nu}_e$ can induce IBD interactions, though their flux in the range of interest (< 10 MeV) is negligible compared to reactor antineutrinos [108]. Otherwise, (anti)neutrinos can undergo various charged and neutral current interactions with nuclei (Z, A) in the detector medium, some of which can produce neutrons via

$$\begin{aligned} \overset{(-)}{\nu} + (Z, A) &\rightarrow \overset{(-)}{\nu} + (Z, A-1) + n, \\ \overset{(-)}{\nu} + (Z, A) &\rightarrow l^{\pm} + (Z \mp 1, A-1) + n. \end{aligned} \tag{6.6}$$

Now, any background producing neutrons at a high enough rate can become a problem for the antineutrino analysis, since two sufficiently coincident neutron captures can pass all the analysis cuts. Thankfully, studies performed in the water phase [22] and then extended to the partial-fill phase [109, 110, 2] using the GENIE neutrino Monte-Carlo generator showed negligible contamination of the antineutrino analysis from this source. Since the full-fill phase is essentially a roughly doubly-scaled up version of partial-fill, these results carry over to the present analysis.

Muons

Cosmic rays can lead to muons passing through or near the detector in two ways: either producing them directly as part of the particle shower, or indirectly via atmospheric neutrinos interacting with nuclei in the nearby rock to produce more muons [111]. The total rate of muons passing through the SNO+ detector is much smaller than most other experiments, being greatly reduced by the significant rock overburden. It was measured both by the SNO experiment [111], and in the SNO+ water phase [112] to be approximately 3 muons per hour.

These high energy cosmic muons can lead to IBD-like concident event pairs, by causing nuclei in the medium to disintegrate, producing neutrons and unstable isotopes [113]. This disintegration and subsequent activity is called muon spallation.

Long-Lived Nuclei Long-lived Spallation daughter nuclei can emit neutrons when they decay some time after the muon has passed. Many possible nuclei can be considered, though those that undergo a $\beta + n$ decay are obviously the most troublesome by once again mimicking the IBD signature. The KamLAND and Borexino collaborations determined that only two isotopes are of primary concern for their antineutrino analyses [114, 115]:

$$\begin{aligned} {}^{9}\text{Li} &\rightarrow 2\alpha + n + \overline{\nu}_e + e^- \\ {}^{8}\text{He} &\rightarrow {}^{7}\text{Li} + n + \overline{\nu}_e + e^- \end{aligned} \tag{6.7}$$

^{9}Li and ^{8}He decay with half-lives of 0.18 and 0.12 seconds respectively, and with Q-values of 13.6 and 10.7 MeV. Their $\beta + n$ decays shown above have branching ratios of 51% ans 15% respectively, and Borexino concluded that between these two, ^{9}Li decays dominate [115]. All other possible muon spallation daughters were considered here [113, 116], either not decaying via $\beta + n$, or having a much lower yield than ^{9}Li and ^{8}He. The one exception is ^{17}N decay, with a half-life of 4.173s [116], though it must be produced via the ^{18}O(n, n+p)^{17}N interaction, and ^{18}O is exceedingly rare in liquid scintillator. This can however be produced in more significant quantities outside the AV, in the UPW shielding.

Neutrons The neutrons inside the AV will capture on hydrogen, releasing the usual 2.2 MeV γ. Therefore, if at least two neutrons are released together, these can mimic the IBD signature. If a neutron has high enough energy (fast neutron) to induce proton or ^{12}C recoils before capturing, it can likewise resemble an IBD event pair, just as for the (α, n) events discussed below. Fast neutrons produced outside the AV may also travel far enough into the detector to be tagged as an event pair [113].

Background Reduction

Two simple but effective cuts are introduced to reduce all atmospheric backgrounds to negligible levels.

Muon Veto Muons passing inside the AV produce a lot of light, and so are easily identifiable high N_{hit} events. Any event with $N_{\text{hit}} > 3000$ is thus tagged as a muon. This corresponds to a roughly 10 MeV reconstructed energy, and so is above the energies of interest in this analysis. The tagging efficiency of muons passing inside the AV was found to be negligibly different from 100% [2].

To remove all the muon spallation following such a tagged event, any events in the following 20s are simply cut. This is to allow the longest-lived concerning isotope ^{17}N around 5 half-lives to decay away. The low muon rate at SNO+ allows such a conservative approach with the loss of only around 7% of livetime. The combination of a low muon rate, high tagging efficiency inside the AV, and a long veto window renders muons a negligible background in this analysis; with the exception of fast neutrons entering the AV from un-tagged muons that passed through the external water. The Multipicity cut described next helps shore up this last possibility, however.

External Muon Veto Meanwhile, muons passing through the water in the cavity outside the AV also produce light that can be picked by the OWLs (outward looking PMTs). Consequently, any event with over 3 OWL hits is tagged, and a 10 μs veto applied. This is to remove fast neutrons that can enter the detector as a result of muon spallation. The tagging efficiency of this method is not as high, due to the lower concentration of OWLs, and the lower light yield of water.

Multiplicity Cut If any event with reconstructed energy over 0.4 MeV takes place between or within ±1 ms of either events of a tagged IBD pair, and less than 2 m away from either, the IBD pair is cut. This eliminates events producing more than two neutrons that may have passed previous cuts, and the possibility of two IBD event pairs overlapping. The low 0.4 MeV energy cut is designed to pick up on any proton recoil events caused by fast neutrons, like those that may enter the AV from external muon spallation.

With all these cuts combined, the atmospheric backgrounds are rendered negligible for this analysis [3, 110], though their individual efficiencies have not been studied in detail.

6.3.2 Geo-neutrinos

Spectra

Geo-neutrinos are the neutrinos emitted by naturally occurring radioactive beta decays inside the Earth. These are thought to occur inside the crust and the mantle, producing the majority of the Earth's internal heat [117]. The four decay chains (^{238}U, ^{235}U, ^{232}Th, ^{40}K) are responsible for next to all this radiogenic heat, and their associated byproducts: electron antineutrinos (and to a lesser extent, neutrinos) [118]. The induced antineutrino spectra are shown in figure 6.4a, though the expected fluxes of these are subject to very large uncertainties, based on which Earth model one uses.

The ^{238}U and ^{232}Th spectra reach energies above the 1.806 MeV IBD threshold, and are thus able to produce IBDs entirely indistinguishable from reactor antineutrino induced ones. The two geo-neutrino IBD spectra are shown in figure 6.4b, only overlapping with the lowest energy reactor antineutrinos. The measurement of this flux is thus an interesting physics goal in its own right, and is effectively measured simultaneously with the present reactor antineutrino analysis, by necessity. No cuts can eliminate this background without also removing the low energy reactor IBDs.

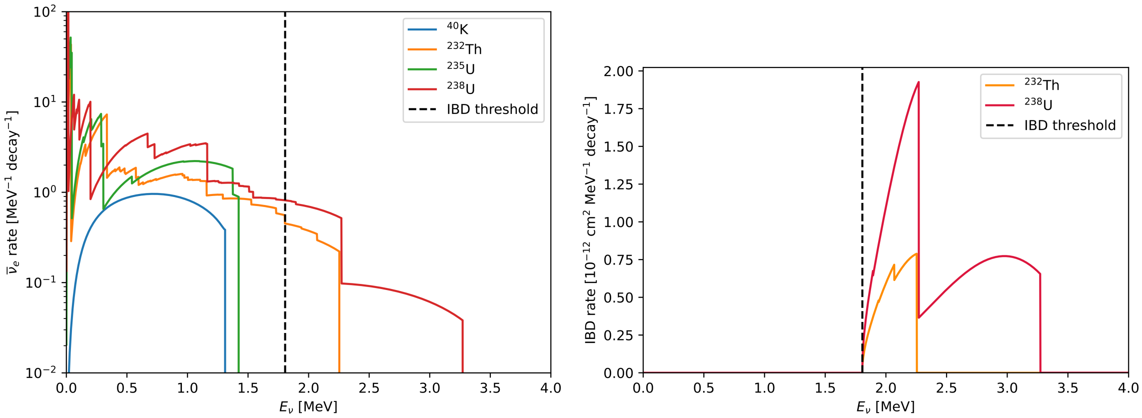

(a) Antineutrino spectra produced by the four primary radioactive decay chains in the Earth, normalised to one full decay through the whole chain. Data taken from Ref. [119, 120].

(b) IBD cross-section applied to the geo-neutrino spectra on the left, using the total cross-section formula and the total number of target protons described in the Predictions chapter.

Figure 6.4: Geo-neutrino spectra with and without the IBD cross-section.

The geo-neutrino flux has been measured by two other experiments thus far: Borexino [121] in Italy and KamLAND [122] in Japan, with tentative first measurements from SNO+ coming out since 2023 [118] and concurrently with this thesis. SNO+ is thus producing the third such independent measurement, and the first in the Americas. These measurements are generally reported in terrestrial neutrino units (TNU), which are defined as the number of IBD interactions per year, per 10^{32} protons, for a fully efficient detector. Bulk-Silicate Earth (BSE) models are a generalised framework to describe the Earth's composition, with the TNU predictions of three such BSE models shown in table 6.1. Meanwhile, a rough summary of current measurements is shown in figure 6.5.

Table 6.1: Geo-neutrino flux predictions in TNU from three types of BSE models, with different heat (Q) productions: either low, mid or high. These are taken from [118], which used numbers provided by Ondřej Šrámek, following the same methods as Ref. [123].

	Low-Q	Mid-Q	High-Q
U	29.72 ± 4.69	34.11 ± 5.04	41.54 ± 4.95
Th	8.21 ± 0.65	9.53 ± 0.77	11.51 ± 0.81

Because of this broad spectrum of models, the total geo-neutrino flux is largely considered unconstrained in this analysis, though the Mid-Q model is taken as a starting assumption. However, a constraint can be applied to the ratio of the two spectra: The ratio of the Uranium to Thorium chain fluxes is [118]

$$R_{\mathrm{U/Th}} = 3.7 \pm 1.3. \tag{6.8}$$

Converting from TNU to an expected number of events at SNO+ (and vice-versa) depends on the detection efficiency of these fluxes, and so is saved for the end of this chapter when all cuts and tagging efficiencies have been discussed.

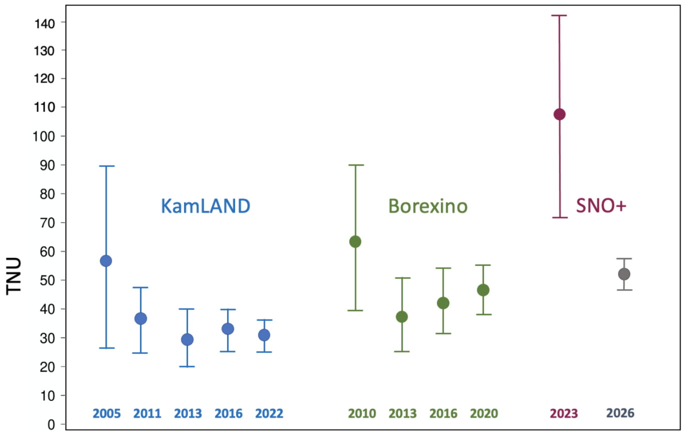

Figure 6.5: Illustration of the measured geo-neutrino fluxes in the world over time. The plot and the SNO+ 2023 result are from Ref. [118], with the grey point being a prediction with 2 years of data and assuming the Mid-Q model.

Oscillation

All the measurements and predictions above have neutrino oscillation built-in. When simply measuring the geo-neutrino flux at SNO+, this is exactly what one wants. However, since this analysis is an oscillation fit, its response to varying oscillation parameters must be understood. One can of course integrate over the full oscillation probability formula, and the local geology (which is a rather unique impact crater around SNO+), as was done in detail in [118]. However, this would be highly computationally demanding to re-compute countless times as is required in an oscillation analysis.

Instead, with geo-neutrinos originating from a continuous distribution spanning hundreds and even thousands of kilometres, instead of from point sources, the oscillation is completely averaged out. As such, formula (5.4) discussed in the Oscillation chapter is very applicable here

$$\langle P_{\overline{\nu}_e \to \overline{\nu}_e} \rangle = s_{13}^4 + c_{13}^4 \left[1 - \frac{1}{2} \sin^2(2\theta_{12}) \right]. \tag{6.9}$$

Oscillation therefore acts as an overall scaling factor, without introducing any appreciable spectral changes. Any deviations from this behaviour are orders of magnitude smaller than the large flux uncertainties previously discussed. This includes potential matter effects, which were found to introduce deviations of less than a percent by the KamLAND experiment [120]. More exact quantification of the impact of the approximations on the spectral shapes seen at SNO+ is currently underway.

One can therefore either fit the oscillated geo-neutrino flux independently from oscillation, or apply oscillation to the un-oscillated geo-neutrino flux simultaneously during the fit, using this scaling formula. The latter option was selected for this analysis, simply for consistency.

6.3.3 (α, n) Events

The major troublesome background of this analysis is caused by α particles capturing on nuclei in the detector, and releasing a neutron which later captures on a hydrogen atom. This broad category is called (α, n) events. The dominant source of α particles inside the AV is from the decay of ^{210}Po, as was described in section 6.2.1. α particles are emitted from the AV surface, while a lower rate is emitted throughout the internal volume.

These can then capture on various particles, such as $^{14/15}$N, $^{17/18}$O or ^{13}C [2]. All possible sources of neutrons were considered by V. Lozza [124], including (α, n), though it turns out that only two types need be considered. Capture on N can be ignored due its negligible presence both in the AV and the liquid scintillator. Oxygen however makes up roughly 37% of the AV mass, and both $^{17/18}$O isotopes are stable. The natural abundance of ^{18}O is over 4 times higher than ^{17}O [125], so only the former was considered. Meanwhile, the liquid scintillator is made up of carbon chains, and so the stable ^{13}C must also be accounted for.

The two noteworthy interactions are thus "surface" ^{18}O(α, n)^{21}Ne in the AV, and "internal" ^{13}C(α, n)^{16}O in the liquid scintillator, or more explicitly

$$\alpha +^{18}\mathrm{O} \;\rightarrow\; {}^{21}\mathrm{Ne} + \mathrm{n}, \qquad \alpha +^{13}\mathrm{C} \;\rightarrow\; {}^{16}\mathrm{O} + \mathrm{n}. \tag{6.10}$$

The former interaction leads to deexcitation γ's of up to 2.1 MeV, while the latter up to 6.1 MeV [8]. In either case the events lead to a prompt event from the interaction's immediate aftermath, followed later by a neutron capture, thus mimicking the IBD signature.

A 5.7 m fiducial-volume (FV) cut was found to reduce the higher rate surface ^{18}O(α, n)^{21}Ne interactions to negligible levels [2]. Therefore, (α, n) refers only to internal ^{13}C(α, n)^{16}O events through the rest of the text, and a 5.7 m cut is imposed throughout. The total rate and prompt energy spectra are constructed through various steps in the following sub-sections.

^{210}Po Rate

^{210}Po decays to ^{206}Pb with a Q-value of 5.407 MeV. It releases a 5.304 MeV α particle, or in rare cases a 4.517 MeV α and a 0.803 MeV γ (0.0012%) [126]. The latter decay branch is ignored due to its rarity. The α energy is quenched down to 0.4 MeV, producing an energy peak that can be fitted, allowing the decay rate to be found. This was performed by S. Ricetto for different FV cuts, and the rates are shown in figure 6.6.

The average rate over this period, for the largest 5 m FV cut, is 40.76 mHz/m^3. Rates computed from smaller FV cuts all smaller by 1.4% or less, and so no strong radial dependence is found within 5 m. Any uncertainty arising from this measurement is negligible compared to the cross-section concerns discussed next.

α to (α, n) Conversion

Thankfully for the antineutrino analysis, only a small fraction of these α decays lead to (α, n) events. This yield can be calculated via

$$Y_{(\alpha,\,\mathrm{n})} = n_{13C} \int_0^{E_\alpha} dE_\alpha \frac{\sigma\left(E_\alpha\right)}{dE_\alpha/dx}, \tag{6.11}$$

where $E_\alpha = 5.304$ MeV is the emitted α energy, $n_{^{13}C}$ is the ^{13}C density in the liquid scintillator, and dE_α/dx is the α's stopping power in the medium. $\sigma\left(E_\alpha\right)$ is the

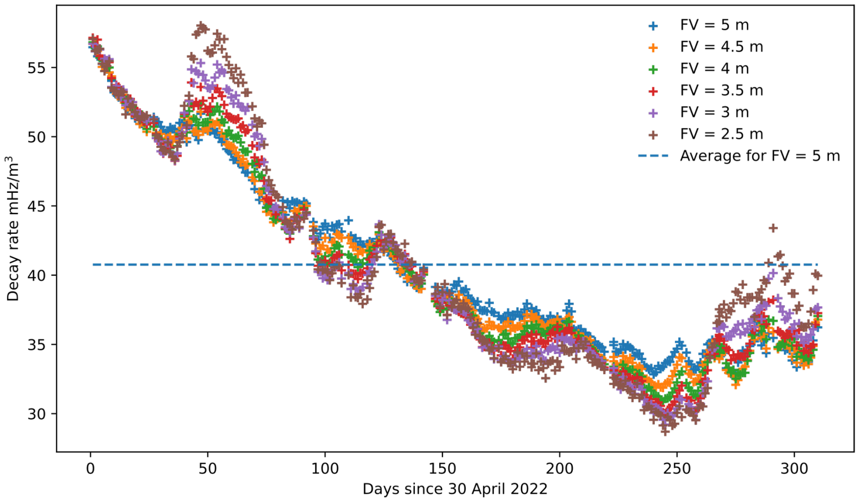

Figure 6.6: Measured ^{210}Po rate between 30/04/2022 and 05/03/2023, with various FV cuts [127].

^{13}C(α, n)^{16}O cross-section itself, which still carries large uncertainties and conflicts between models and measurements. The JENDL/AN-2005 libraries were used for this analysis [128], which are cross-section calculations based on measurements by K.K. Sekharan et al. (1967) [129], J.K. Bair et al. (1973) [130], and others. These are shown in figure 6.7, along with the more recent Harissopulos et al. measurements [131]. The stopping power was calculated using SRIM [132], and also shown in figure 6.7.

The ^{13}C density can be calculated similarly to the H density in the IBD flux calculation, in section 4.4.2. Using the mass fraction information for carbon, and a naturally abundant average carbon atom mass of 12.0107 amu [134], one finds a carbon density of $n_C = 3.79 \times 10^{22}\text{cm}^{-3} \pm 0.5\%$ at 15°C (uncertainty from temperature dependence). The molar fraction of ^{13}C is 1.07% [134], which finally modifies this to

$$n_{13C} = 4.05 \times 10^{20}\text{cm}^{-3} \pm 0.5\%. \tag{6.12}$$

The (α, n) yield was previously computed by V. Lozza [124], and then updated to [135]

$$Y_{(\alpha,\text{ n})} = 6.11 \times 10^{-8}. \tag{6.13}$$

Daya Bay carried out a similar calculation, finding a yield of $(5.8 \pm 0.4) \times 10^{-8}$ in a Gd-loaded LAB liquid scintillator, similar to that of SNO+ [136]. It has a carbon mass fraction of 87.7%, almost identical to the SNO+ scintillator (see table 4.4), though it then computes the ^{13}C mass density by multiplying the carbon mass density by the ^{13}C molar fraction (1.1%), rather than the mass fraction, yielding 0.0083 g/cm^3. This would lead to a ^{13}C density of $3.85 \times 10^{20}\text{cm}^{-3}$ which, accounting for the slightly different ^{13}C fractions used (1.1% vs 1.07%), is about 8% lower than the SNO+ value. Correcting Daya Bay's (α, n) yield for this brings it into agreement with the value obtained here (6.13). Discussion of uncertainty is reserved for a following section.

Bringing the (α, n) yield together with the aforementioned average ^{210}Po decay rate per unit volume (in the 5 m FV), the total (α, n) rate in the 5.7 m FV is

$$r_{(\alpha,\text{ n})} = 1.93\mu\text{Hz}, \tag{6.14}$$

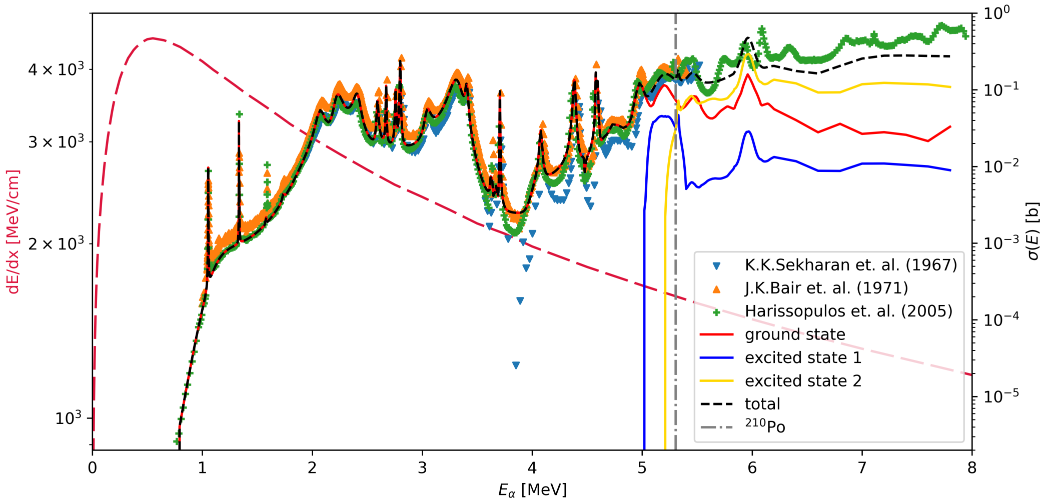

Figure 6.7: (α, n) total and partial cross-section calculations from the JENDL/AN-2005 libraries [128], along with total cross-section measurements [129, 130, 131], are shown on right y-axis. All these were downloaded from Ref. [133], in the JENDL/2005 and EXFOR branches. The α energy from the ^{210}Po decay is shown as a vertical line. The left y-axis shows the α stopping power in liquid scintillator, calculated using SRIM [132].

which translates to 61 expected (α, n) events per year, ignoring detection efficiency.

The angular correlation between the α and the neutron is taken into account only in the ground state case, modelled with Legendre polynomials, just as in KamLAND [120]. The impact of this angular dependence on simulated prompt energy spectra was studied by the Daya Bay and KamLAND experiments [136, 120], and was found to be small, due primarily to the neutron's random walk.

Prompt Spectrum

As shown in figure 6.7, the ^{16}O can be produced in either its ground state, or one of its two excited states reachable by the ^{210}P's α energy. The neutron is released with between 2.2 and 7.5 MeV for the ground state case, and lower energies when the ^{16}O is excited, as shown in figure 6.8. This leads to three possible prompt signals, which do not correspond one-to-one with the three excitation states:

- Proton recoil (PR): this is the dominant process. With the ^{16}O produced in the ground state, the neutron takes away almost all the energy – a so-called fast neutron – and then elastically scatters off multiple hydrogen atoms in the medium, losing energy over time. Simulations with GEANT4 reveal an average of 7 recoils, during the first few ns of the neutron's random walk [2]. These recoiling protons then go on produce quenched scintillation light, which sums to events from around 0.4 to 3.5 MeV.
- ^{12}C scattering: the same fast neutron can instead inelastically scatter on a ^{12}C atom, raising it to an excited state, which rapidly decays by releasing a 4.4 MeV γ [137]. The fraction of ^{12}C scatter versus PR events is estimated by simulation, using cross-sections from [138].

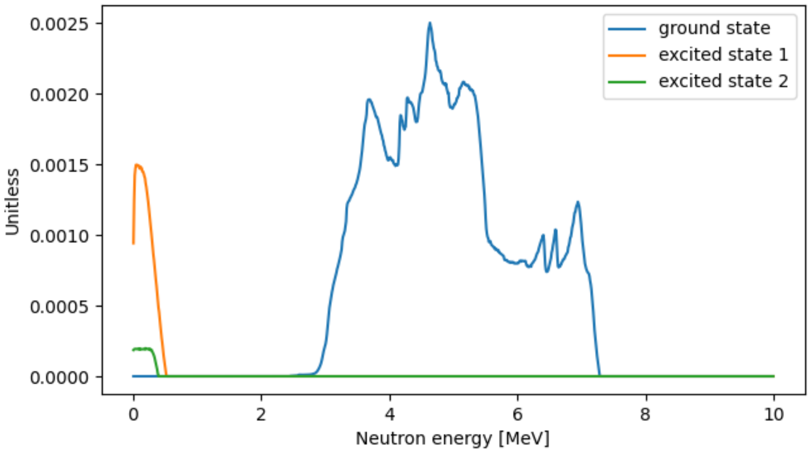

Figure 6.8: Emitted (α, n) neutron energy spectra in SNO+ scintillator, for ^{16}O produced in its ground or first two excited states. Computed analytically [124].

- ^{16}O deexcitation: the resulting oxygen atom is produced in either its first (6.049 MeV) or second (6.13 MeV) excited states, with α energy thresholds of 5.014 and 5.119 MeV respectively [128]. The former de-excites by releasing a e^+e^- pair, the latter via a γ. Due to energy smearing this forms a single peak around 6 MeV, so that both processes are indistinguishable.

Some energy from the original α contributes too, though its fraction is very low, with a mean α energy of 0.2 MeV after quenching seen in the KamLAND experiment [120]. Otherwise, roughly 9% of (α, n) prompt events are ^{16}O deexcitation, 2.3% ^{12}C scattering and the rest PR, which can be seen as distinct peaks in the energy spectrum of figure 6.9. There is however a great deal of uncertainty in these fractions, and indeed in the overall yield.

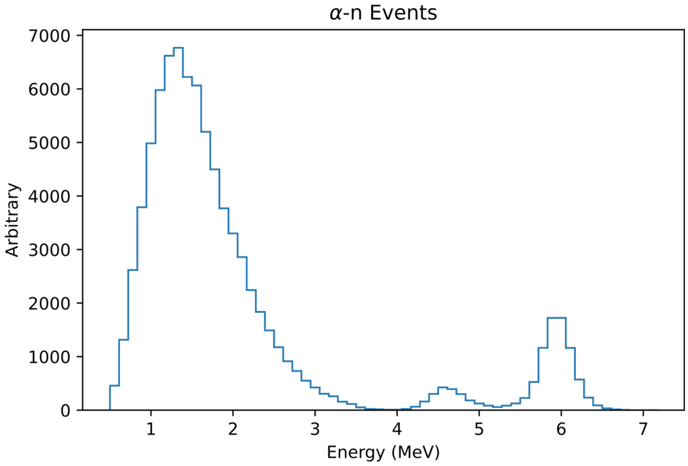

Figure 6.9: MC generated (α, n) prompt energy spectrum, with the three prompt event processes clearly visible as distinct peaks. In order from left to right: proton recoil (PR), ^{12}C scattering and ^{16}O deexcitation.

Uncertainties

Different total cross-section measurements show significant disagreement, such as between [131] and [130], among others. Even more troublesome, the ground state is based on measurements of the inverse ^{16}O(n, α)^{13}C process, while the excited states are simply calculated [128], with scarce experimental verification. The lack of certainty here is such that the IAEA's nuclear data section (NDS) put together a summary report laying out these issues, and the highest priority measurements [139]. In addition, the quenching experienced by protons in liquid scintillator is often hard to independently calibrate, so there is uncertainty in the PR's reconstructed energy.

In order to address these concerns, the KamLAND experiment assigned conservative uncertainties to the (α, n) prompt energy PDFs used in their spectral and oscillation analyses. The ground state PDFs (PR and ^{12}C scattering) were given a joint 30% normalisation uncertainty, while the excited state PDF (^{16}O deexcitation) is left entirely unconstrained. An extra 10% uncertainty in the PR's energy scale is also implemented, to account for the quenching uncertainty [120]. The present analysis will largely replicate this procedure, described in detail in chapter 8.

6.4 Tagging Efficiencies and Rates

Now that all the players are accounted for – the backgrounds, and the cuts used to reduce some of them – the effectiveness of these cuts on the backgrounds and the IBD signal can be determined. All the cuts are summarised in table 6.2, but first a distinction must be made: any cuts involving timing vetoes (the muon induced vetoes) are treated as reducing the live-time of the detector, as if the detector were off during these periods. They therefore do not appear in tagging efficiency calculations, but must be reflected in reported live-times. These reduce the total live-time of the dataset from 145.2 (sum of all the runs) to 134.5 days, explaining the value used throughout this text. Any other cuts conversely are only ever treated as tagging efficiency, without impacting the detector live-time. The distinction is made for ease of implementation in calculations, since everything but the vetoes is computed using MC simulations. This section concerns itself only with the latter case.

The efficiency of the FV cut (R cut) has a little subtlety, in that is depends on how one defines it. The full detector volume in which RAT simulates events – called the "logical" volume – is 920.23 m^3, slightly larger than a 6 m sphere by including some of the neck. However, a 5.7 m FV cut is always imposed, and was in fact used for all background rate calculations in this chapter, as well as for the reactor IBDs (see section 4.4.2). This allowed a reduced uncertainty in the number of target protons, by not needing to estimate the full detector volume. Therefore, only events simulated within 5.7 m shall be considered for these efficiency calculations. Any residual R cut efficiency deviations from 100% are thus purely from radial bias and resolution in position reconstruction.

For the multiplicity cut, let us assume the event rates are constant, and the worst case largest multiplicity window of $T_{\mathcal{M}} = 2.8$ ms. The highest rate event is from reactor IBDs, at a rate after all other cuts of 4.20×10^{-6} Hz. The number of events falling in this window is Poisson distributed, with an expectation value of $\lambda = T_{\mathcal{M}} \cdot r_{\text{IBD}} = 1.2 \times 10^{-8}$. The probability of at least two events falling in this

Table 6.2: Summary of cuts applied throughout this analysis. E and R are the reconstructed energy and radial positions of the event respectively. ΔR and Δt are, in order, the spacial and temporal separation of the prospective prompt and delayed events. Valid signifies whether reconstruction was successful or not, while other quantities are described in sections above.

Cut name	prompt cut		delayed cut		coincidence cut	
	min	max	min	max	min	max
E (MeV)	0.9	8	1.85	2.4		
R (m)		5.7		5.7		
ΔR (m)						1.5
Δt (μs)					0.4	800
Valid	True					
Multiplicity	$\mathcal{M} = 1$ in $[t_{\text{prompt}} - 1\text{ms},\ t_{\text{delayed}} + 1\text{ms}]$					
Muon Vetos						
Internal	20 s after $N_{\text{hit}} > 3000$					
External	10 μs after OWL $N_{\text{hit}} > 3$					

window is thus approximately zero, and so the multiplicity cut does not meaningfully affect the tagging efficiency of any signal or background except what it is designed to suppress.

The efficiencies of these cuts on the reactor IBD signal are shown in table 6.3, along with any backgrounds that were not rendered negligible. These are all computed using MC simulations, except for the data-driven accidentals background, for which the numbers come from the arguments in section 6.2.2 above. Notice the 96% FV cut efficiency, which likely comes from reconstruction's roughly 10 cm position resolution at around 5.7 m [67]: the volume of the shell from 5.6 to 5.7 m is roughly 96% of that from 5.7 to 5.8 m, meaning events mis-reconstructed ±10 cm about 5.7m have more chance to be reconstructed outside 5.7 m than inside by a ratio of about 96%. This resolution effect, as well as any potential minor radial bias, is assumed to be the same for the reconstruction of both MC simulations and data, and so is not included in any systematic uncertainties at this stage.

The expected event rates after all these cuts are shown in table 6.4. The unoscillated geo-ν flux $N_{\text{geo-}\nu}$ is computed using the Mid-Q TNU flux Φ from table 6.1, the $R = 3.7$ U/Th radio, and the tagging efficiencies ϵ_{U} and ϵ_{Th} from table 6.3, via

$$\Phi[\text{TNU}] = \frac{\langle P_{ee}\rangle N_{\text{geo-}\nu}}{R+1}\left(\frac{R}{\epsilon_{\text{U}}} + \frac{1}{\epsilon_{\text{Th}}}\right)\frac{10^{32}}{N_H}\frac{1\text{yr}}{T}. \tag{6.15}$$

$N_H = 4.849 \times 10^{31}$ is the number of protons in the target volume, as computed in section 4.4.2, while $T = 134.5$ days is the livetime, and $\langle P_{ee}\rangle = 0.54$ is the average survival probability. The expected number of oscillated reactor and geo-ν IBDs with PDG oscillation parameters are also shown in the table, producing a total of 50.2 expected events. The resulting expected prompt energy spectrum is shown in figure 6.12.

Table 6.3: Cut efficiencies (%) for the various cuts applied to the signal and non-negligible backgrounds in this analysis. These are all computed using (un-oscillated) MC simulations, except for the accidentals background, which is added as a rough comparison, and assumes a starting "un-cut" rate given from all reconstructed events within the 5.7 m FV, and between 0.5 and 9 MeV. Note that all the individual cut efficiencies are only approximate, and depend on their ordering. The final cut efficiencies themselves are accurate however, with statistical uncertainties of 0.1% or less. Further note that all these efficiencies are for events simulated within the 5.7 m FV, and the Valid cut efficiency includes the detector and reconstruction efficiencies, and a $E > 0$ condition to deal with non-sensical energy reconstruction.

<table>
<tr><th></th><th>Cuts</th><th>reactor-ν</th><th>geo-ν (U)</th><th>geo-ν (Th)</th><th>(α, n)</th><th>accidentals</th></tr>
<tr><td rowspan="3">delayed</td><td>Valid</td><td colspan="4">⟵ 98 ⟶</td><td>-</td></tr>
<tr><td>R</td><td colspan="4">⟵ 96 ⟶</td><td>-</td></tr>
<tr><td>E</td><td colspan="4">⟵ 96 ⟶</td><td>0.0715</td></tr>
<tr><td rowspan="3">prompt</td><td>Valid</td><td>99</td><td>98.5</td><td>98</td><td>96.5</td><td>-</td></tr>
<tr><td>R</td><td colspan="4">⟵ 96 ⟶</td><td>-</td></tr>
<tr><td>E</td><td>99.9</td><td>99</td><td>98</td><td>92</td><td>1.46</td></tr>
<tr><td></td><td>Δt</td><td colspan="4">⟵ 97 ⟶</td><td>0.680</td></tr>
<tr><td></td><td>ΔR</td><td colspan="4">⟵ 99 ⟶</td><td>8.5</td></tr>
<tr><td></td><td>Total</td><td>82.9</td><td>81.9</td><td>80.8</td><td>73.3</td><td>6.03×10^{-9}</td></tr>
</table>

Table 6.4: Average expected event rate after tagging efficiency, assuming a geo-ν U/Th ratio of 3.7 as described previously. The expected number of events for 134.5 days is also shown, with and without neutrino oscillation.

	reactor-ν	geo-ν	(α, n)	accidentals
Rate (Hz)	4.20×10^{-6}	1.01×10^{-7}	1.41×10^{-6}	5.12×10^{-8}
Events (no oscillation)	48.8	11.8	16.4	0.59
Events (with oscillation)	26.8	6.36	16.4	0.59

6.5 Tagged Data

These same cuts were applied to to the real 134.5 day dataset, and produced 57 tagged event pairs, within the Poissonian uncertainty about 50.2 ± 7.1. Their locations are displayed in figure 6.10a, exhibiting a roughly uniform distribution throughout the FV, as expected for IBDs and approximately for (α, n) and accidentals. The prompt and delayed events' average x, y and z positions range from -0.09 to 0.37 m, with standard deviations of roughly 2.5 m. Since almost all events are expected to be coincident neutron capture events – IBDs or (α, n) – the tagged events should follow the IBD's Δt and ΔR distributions, while their delayed event energies should also match up. The former two are presented in figure 6.11, and the latter in figure 6.10b. These all match up with expectations, given the limited statistics[1].

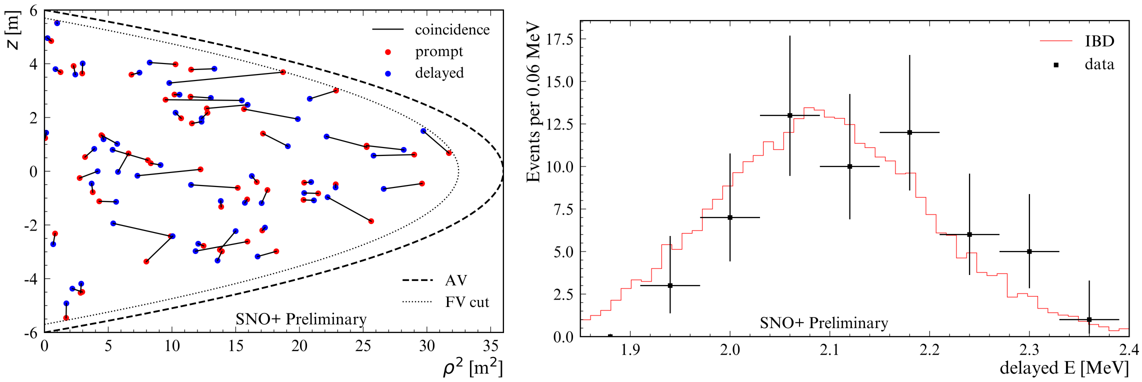

(a) Event pairs displayed according to their reconstructed z and $\rho^2 = x^2 + y^2$ positions in the detector.

(b) Reconstructed energy of delayed events compared to simulated IBD delayed events. The IBD normalisation is made to match the data.

Figure 6.10: Tagged prompt and delayed events from data taken between 30/04/2022 and 05/03/2023.

Meanwhile, the quantity of greatest interest – the prompt event energy distribution – of these tagged pairs is shown in figure 6.12, compared to the expected one from all the above calculations. This spectrum is more complex than all the others mentioned here, depending on the contributions from the reactor IBD signal as well as its relevant backgrounds, and their relative fluxes. PDG values of oscillation parameters were also assumed for this calculation, with which the data and simulations appear to be compatible. It is clear that despite the as yet small sample size, its distribution in prompt energy is sensitive to these details, in particular to the contribution of neutrino oscillation.

6.6 Summary

The final predictions of the fluxes and distributions of the reactor IBD signal and its backgrounds were established, accounting for detection and background reduction tagging efficiencies. These allow for a prediction of approximately half of the tagged events originating from reactor antineutrinos, along with a likely detectable contribution from geo-neutrinos. The 57 tagged events from the analysed dataset were

---

[1] A note on the data error bars used in this thesis is in appendix A.2

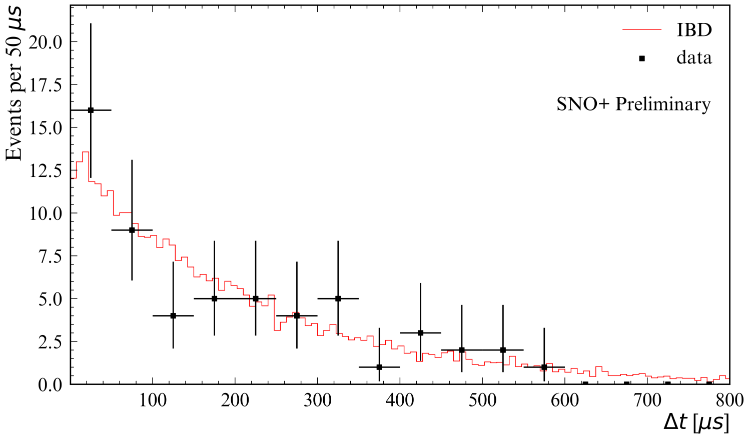

(a) Temporal coincidences.

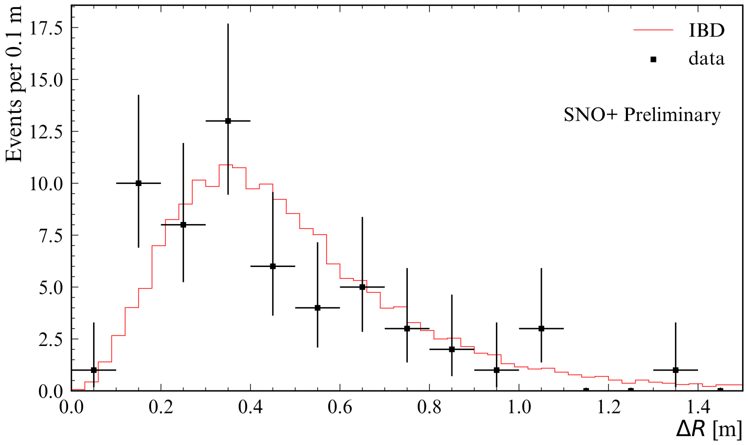

(b) Spacial coincidences.

Figure 6.11: Coincidence quantities between tagged prompt and delayed events from data taken between 30/04/2022 and 05/03/2023, compared to expected distributions from MC simulated IBDs. The IBD normalisations are made to match that of the data.

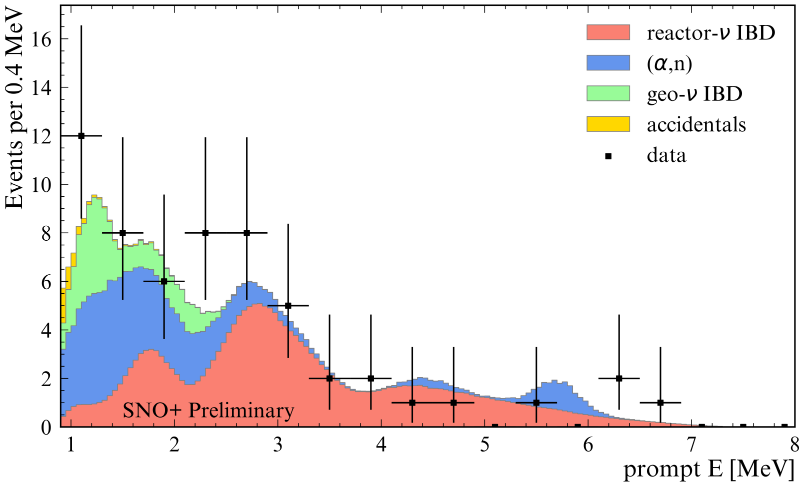

Figure 6.12: Reconstructed energies of tagged prompt events from data taken between 30/04/2022 and 05/03/2023, compared to the expected total prompt energy spectrum from all signal and backgrounds in the same period, from MC simulations. These use the predictions from table 6.4 and PDG values of oscillation parameters [6].

found to match well with all the predictions developed across this and the previous chapters, within statistical uncertainties. Meanwhile, systematic uncertainties from all contributions were discussed, and are summarised below for convenience.

The uncertainty in the reactor antineutrino IBD flux was already covered in previous chapters, being around 3%. The geo-ν flux is left unconstrained, while the U/Th ratio is assigned a 35% uncertainty. The (α, n) flux is split into its three prompt signals, with a joint 30% normalisation uncertainty for the two first (ground state products: proton recoil and ^{12}C scattering), and a 100% uncertainty in the third peak's (^{16}O de-excitation) normalisation. The proton recoil distribution is also given an additional energy scaling uncertainty, though this will be covered in chapter 8. Finally, the accidental background has a 40% normalisation uncertainty stemming from the ΔR cut, as previously mentioned, but no energy systematics due to its PDF being data-driven. It is also assumed that the detection and tagging efficiency contribute negligible uncertainty to any of these signals and backgrounds. Changes in detector state are taken into account in simulations, while systematic uncertainty in the energy reconstruction of prompt events is taken into account independently in chapter 8. The impact of energy systematics on the delayed energy cut efficiency is not accounted for in this work however.

In terms of potential improvements to background reduction, the accidental coincidence rate is already almost negligible, though small improvements can be made, as in [3]. Otherwise, the geo-ν flux cannot be separated from the reactor-ν flux by its nature. This leaves the (α, n) background as the most troublesome remaining background to eliminate, which is the topic of the next chapter.

Chapter 7

(α, n)-IBD Classifier

> "And he that breaks a thing to find out what it is has left the path of wisdom."
> Gandalf – *The Lord of the Rings*, J.R.R. Tolkien

7.1 Introduction

After the coincidence tagging introduced in the previous chapter, the largest remaining background to the reactor antineutrino signal is (α, n) events. This background is quite troublesome for a few reasons, visible in the spectrum of figure 6.12. First its various peaks can somewhat interfere with the recognisable spectral shape induced in the reactor IBD spectrum by Δm^2_{21}. Secondly, the proton recoil (PR) peak – the largest by far – completely overlaps with the geo-neutrino spectrum, rendering these two distributions largely degenerate in a spectral analysis. Thirdly, it is plagued with uncertainties, as discussed in chapter 6. These systematic uncertainties reduce the sensitivity of the oscillation analysis, and have an even larger impact on geo-ν flux measurements. Only a qualitative explanation is given here, but all this will be borne out by the oscillation analysis results in the following chapter.

The saving grace here is that proton recoils have a fundamentally different event structure to the electron-positron annihilation of prompt IBD events, and indeed most events detected at SNO+. The high energy neutron from the (α, n) recoils against protons in the medium, producing an event slightly more spread out in time. While large liquid scintillator detectors are insensitive to tracks and directionality, they are rather good at extracting timing information. A classifier based on pulse shape discrimination is thus developed to remove most of this background, detailed in what follows. First, the concept of pulse shape discrimination is explored, laying out the information the pulse shape contains and what parameters may affect it. The development of the classifier is then laid out, followed by demonstrations of its performance on various simulated datasets. Finally it is put to task on the real tagged events from the previous chapter.

7.2 Pulse Shape Discrimination

Pulse shape discrimination (PSD) is the practice of classifying particles based on their pulse shapes, which are themselves representations of the temporal structure of an event. For example in a liquid scintillator experiment, this would be the light emission profile over time of an event. The exact way this is obtained at SNO+ is described next, followed by what information it holds.

7.2.1 Time Residuals

One of the only pieces of information available at a fundamental level at SNO+ is when and where a photon was detected by a PMT: the time and location of a PMT hit. These are used to reconstruct the energy, position and time of an event, as discussed in section 3.6.2. Subsequently, the PMT hits are still available for other purposes, and as such are used to obtain pulse shapes. The pulse shape of an event is defined by the distribution of its residual hit times t_{res} from each PMT, given by equation (3.1) in the same section, and repeated here for clarity

$$t_{\text{res}} \equiv t_{\text{hit}} - t_{\text{evt}} - t_{\text{TOF}}. \tag{7.1}$$

The pulse shape of an individual event is not a smooth curve, due to a limited number of PMT hits. However, averaging over a large number of events effectively reproduces a PDF from which they are drawn, as shown in figure 7.1 for (α, n) and reactor IBD events. The same cuts as those presented in table 6.2 were applied to MC simulations, with an additional maximum prompt energy cut of 3.5 MeV, to only look at energies where PR events are present. Note that the pulse shapes shown are not corrected for the different energy spectra of the events summed over, which can introduce an additional normalisation difference, due to the total number of PMT hits (N_{hit}) being roughly proportional to the event energy. However, even ignoring this, the two different shapes are clearly visible: PR events have a longer tail at the expense of a slightly shorter peak.

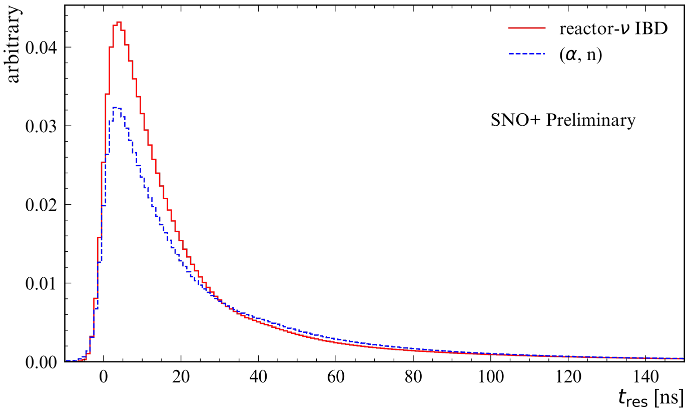

Figure 7.1: Example IBD (red) and α-n (blue) pulse shapes, created by summing over the residual hit times of many simulated events in the 0.9 to 3.5 MeV range.

7.2.2 Scintillator Timing: Proton vs β

Organic scintillators, such as the one employed by SNO+, contain aromatic rings, which produce photons (scintillate) via the excitation and then de-excitation of highly de-localised electrons. It is therefore a type of fluorescence. Any scintillator used in a liquid scintillator experiment also has a Stokes shift, meaning that the photons emitted by it have too little energy to re-excite it. For organic scintillators, this achieved via vibrational energy levels. Due to spin selection rules, there are two scintillator emission components: primary fluorescence and delayed fluorescence. These occur with characteristic timescales from $\mathcal{O}(1)$ ns to $\mathcal{O}(100)$ ns, respectively. Now, excitation via ionisation and recombination tends to lead to a higher fraction of delayed fluorescence, compared to the fluorescence induced by excitation from elastic scattering. As a result, heavy ionising particles such as α's give rise to more delayed fluorescence compared to lighter particles like β's, with protons somewhere in-between. The scintillation pulse from heavier particles thus tends to have a longer tail, which can translate to a longer tail in an event's pulse shape [140, 141, 20]

In RAT, the scintillator emission time PDF – or scintillator timing for short – is parametrised via

$$f(t) = \sum_{i=1}^{n} N_i \frac{e^{-\frac{t}{\tau_i}} - e^{-\frac{t}{\tau_{\text{rise}}}}}{\tau_i - \tau_{\text{rise}}}, \tag{7.2}$$

where N_i and τ_i are the particle-specific decay constants, and τ_{rise} is the scintillation light rise time. Heavier ionising particles with a higher fraction of longer fluorescence times are thus modelled with a longer tail in $f(t)$. $\tau_{\text{rise}} = 0.85$ ns is the same for all particles, while the rest are calibrated by tuning MC simulations to data of known pure samples of the particles of interest. β timing is calibrated using the prompt events of the tagged in-situ ^{214}BiPo event pairs described in section 3.7.2. Meanwhile, α timing takes advantage of tagged ^{210}Po decays, as mentioned in section 6.3.3. The resulting timing constants for the 2.2 g/L PPO loaded full-fill phase used in this analysis are shown in table 7.1. The scintillator emission PDFs thus produced[1] are shown in figure 7.2.

Table 7.1: Scintillator timing constants for β and α particles, used in equation (7.2).

	β timing			α timing			
i	1	2	3	1	2	3	4
τ_i (ns)	5.0	24.46	399.0	4.1	21.0	84.0	197.0
N_i (unitless)	0.656	0.252	0.092	0.523	0.656	0.252	0.092

The proton timing has not been tuned as of the writing of this thesis. Attempts were made to use the externally deployed AmBe source (see section 3.7.3) for this purpose, but issues with simulating neutron propagation into the AV rendered the tuning challenging. The proton timing is expected to lie somewhere between those of the β and α's, for the reasons laid out above. Consequently, the proton timing is set to the same as the α timing, since it was found to fit reasonably well during the partial-fill phase, and better than the β timing [1]. Such a check was possible during the partial-fill phase, since the (α, n) rate was far higher, and so a pure sample could be tagged. This is no longer possible in the full-fill phase. However, the classifier will turn out to be quite insensitive to the small differences in timing expected for protons, as will be shown later in this chapter.

---

[1] Note that the values saved in RAT are confusingly $-\tau_i$, and the formula used is the CDF based off the PDF above, normalised to have its last value equal to 1

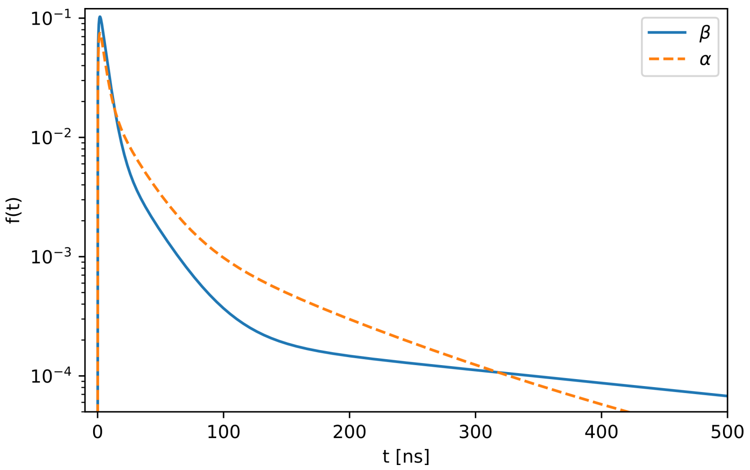

Figure 7.2: Scintillator emission PDFs for α and β particles in 2p2 g/L PPO loaded full-fill optics. Uses constants from table 7.1 substituted into equation (7.2), and then normalised to 1.

7.2.3 Event Structure

Energy Dependence

The timing considerations above affect the relative timing of emitted photons, but not the number of them. The relationship between the light produced by an organic scintillator and the energy loss of a particle travelling through it is well described by Birk's law

$$\frac{dL}{dx} = S\frac{\frac{dE}{dx}}{1 + k_B\frac{dE}{dx}}. \tag{7.3}$$

This is an empirical formula, where L is the light yield, S is the scintillator efficiency, k_B is Birk's constant and x parametrises the particle's path. dE/dx is called the particle's specific energy loss, or stopping power (as was used in equation (6.11)). The value of k_B depends on the particle and scintillator, and was found for β's in LAB using a Compton spectrometer with electrons in the energy range 0.09-3 MeV [142]. The scintillator efficiency was then found by fitting ^{214}BiPo prompt event simulations to data [143]. Birk's constant for α's was similarly found using ^{214}BiPo delayed events [143], while the value for protons was measured with an external experimental setup in a mix of LAB and 2.2 g/L of PPO [144]. All these values are listed in table 7.2, which were used for simulations in this analysis.

Table 7.2: Constants used for different particles in Birk's formula (7.3).

S [photons/MeV]	k_B [mm/MeV]		
All	β	p	α
14000	0.074	0.0708	0.077

The different values of k_B, and the different energy deposition rates (dE/dx) of heavier particles such as α's and protons, lead to an effect known as quenching. Quenching is where particles with the same kinetic energy but higher mass appear to be of lower energy because they have a lower light yield. The exact mechanism is not fully understood, but it is measured, such as the quenched low energy ^{210}Po α decay peak at 0.4 MeV discussed in section 6.3.3.

Nevertheless, one can see that the total light yield is approximately proportional to the energy of a particle, particularly at low energies. This where the rough proportionality between the number of PMT hits and the energy of an event comes from. However, Birk's law also shows that the light yield over time depends on the energy deposition rate, which itself depends on the particle's initial energy. Therefore, not only does the magnitude of the time-residual distribution of an event depend on its energy, but so does the time-residual shape itself. This is further compounded by the interaction-specific energy dependencies below.

Proton Recoils As mentioned in section 6.3.3 of the previous chapter, the (α, n) fast neutron recoils an average of 7 times on protons in the medium over the course of a few ns, as it starts its random walk. It generally deposits less energy with each successive collision, with the exact energy deposition over time having some dependence on the initial neutron energy, though such behaviour has not been formally studied at SNO+ yet. Some of the prompt energy also comes from the instigating α particle, though the amount is low (roughly 0.2 MeV, see section 6.3.3). All this stretches out the event over time, leading to a longer tail in the pulse shape. The way all this responds to different event energies may also be non-trivial, though figure 7.3a shows very little shape-energy correlation, conversely to IBD positron annihilations.

Positron Annihilation Similar arguments can also be used for the IBD prompt events. The event energy comes from two sources: the positron travelling through the medium, and the photons released when it annihilates with an electron. Both the original e^+, and the annihilation γ particles also exhibit energy dependant deposition rates, making the pulse shape correlated with event energy. This correlation can be even larger in the case where the e^+ forms a bound state with an e^- in the medium lasting a few ns, as described below.

When a positron is released, it can sometimes form a bound state with an electron – after having lost enough energy – called Positronium (Ps). The fraction of positrons that undergo this process is both energy and medium dependant [145, 146, 147]. However, out of those that do, a quarter form para-Positronium (p-Ps) with spin 0 (e^+ and e^- spins anti-parallel), while three quarters form ortho-Positronium (o-Ps) with spin 1 (spins aligned), due to combinatorics. The former has a lifetime of 125 ps in a vacuum, making it indistinguishable from the case where Ps is not formed. In fact, the pulse shapes of e^+ and e^-, when corrected for the annihilation energy, are practically identical since they have the same charge magnitude and mass. That is, except for when o-Ps is formed. o-Ps has a lifetime of 142 ns in a vacuum, though it is reduced in a medium by interactions with other electrons, such as "pick-off", "conversion" (spin-flip) or chemical reactions. This reduces the lifetime to roughly 3 ns, which is just enough to distort the pulse shape, particularly for the longer lived cases. This effectively separates the pulses from the e^+ and γ energy depositions, broadening the resulting pulse shape slightly.

The exact values to use for the o-Ps total formation fraction and lifetime are not known, since no dedicated measurement exists for the SNO+ liquid scintillator cocktail. Looking at literature on the topic, the authors in Ref. [148] studied a range of scintillator mixes, and Ref. [149] interestingly investigated the effect of adding Neodynium to the mix, in anticipation of SNO+'s previous plans to use this as a double beta decay source. However, with Tellurium being used instead, Ref. [150] provides the best overall match to the SNO+ target mix: LAB + 2g/L PPO + 5% amine group as a surfactant + various concentrations of Te or Te and water. The only mismatch is their use of an amine group as opposed to bis-MSB. The effect of this is unknown, but from the values in the other papers mentioned above, it is expected to be small, particularly for τ. The Tellurium and water concentrations do not have much impact on τ or f. Meanwhile, theoretical models do not predict much energy dependence on τ, where electric field strength was used simulate different energies [145, 151]. This appears to be in agreement with experimental measurements [145] in the low energy region (higher energies presented much larger backgrounds). The formation fraction is however clearly energy dependent, as mentioned previously. The values introduced into RAT by the author, and thus used in simulations are

$$\tau = 2.7\text{ns}, \qquad f = 0.36. \tag{7.4}$$

The effect of o-Ps is emulated by delaying the γ particle formations, of a fraction f of positrons, by a random number sampled from an exponential decay distribution governed by τ.

No uncertainty has been placed on these values yet, though the impact of including these effects at all is extremely small. Nevertheless, this adds some small extra randomness to the pulse shape, with some additional complicated energy dependence: the initial e^+ peak scales with the positron's initial energy, while the latter annihilation peak does not. All this discussed energy correlation is shown in figure 7.3a.

Radial Position Dependence

The last quantity that has a significant impact on the pulse shape is how far away the event is from the centre of the detector. Due to the changes in refractive index between the liquid scintillator, the AV, and the external water, the detected photons do not necessarily travel in a straight line from the event to the PMT. This is to be expected, as is a certain amount of scattering and re-emission of light even away from these boundaries. However, the geometry of the detector means that these optical effects have a non-trivial dependence on the radial position of an event. This is somewhat taken into account during event reconstruction, though with increasing difficulty as the event approaches the AV edge and the optics become ever more challenging. However, the simple time residual definition (3.1) always assumed a straight line time-of-flight. The discrepancy between this assumption and the reality thus changes with radial position, and as such pulse shapes are warped as a function of this. Note that this is another reason to only look at events within a 5.7 m FV, avoiding the most unreliably reconstructed events and the strangest pulse shapes. Examples of this radial position dependence are shown in figure 7.3b.

7.3 Classifier Selection

With the proton recoil events from (α, n) having longer tails in their pulse shapes on average, compared to the IBD's positron annihilations, a suitable classifier can

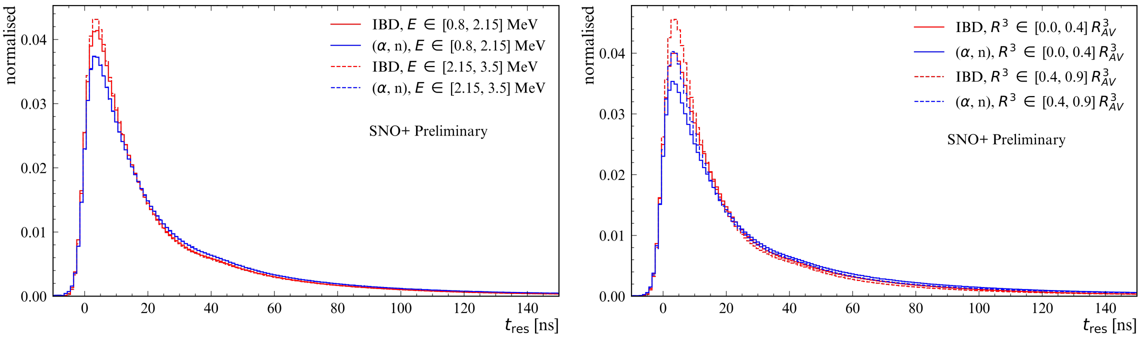

(a) Reconstructed energy dependence. Normalisations also roughly scale with energy.

(b) Reconstructed radial position dependence ($R_{\mathrm{AV}} = 6$ m).

Figure 7.3: IBD and (α, n) prompt event pulses' shape dependence on energy and radial position. Only events with reconstructed energy in the 0.8 to 3.5 MeV range and reconstructed radial position below 5.7 m were considered. Pulse shapes are also always normalised to draw attention to the shape changes.

be constructed to optimally distinguish between them. In particular it should be resilient to the energy and radial position correlations just discussed. The standard log-likelihood classifier is briefly covered for reference, before an in-depth description of the better adapted Fisher discriminant classifier.

7.3.1 Log-Likelihood

For a given PDF $p_A(x)$, the probability density of making a measurement between x_1 and x_1+dx is $dP_A = p^A(x_1)dx$. For N independent measurements x_i, the probability of this outcome is $dP_A = \mathcal{L}_A(x)dx$, where the likelihood $\mathcal{L}$ is defined as

$$\mathcal{L}_A = \prod_{i=1}^{N} p^A(x_i). \tag{7.5}$$

The likelihood ratio $r_{\mathcal{L}} \equiv \mathcal{L}_A/\mathcal{L}_B$ between two hypotheses A and B is then a test of which is most likely to produce a given data set $\{x_i\}$. In practical applications, the logarithm of the ratio is easier to compute: the log-likelihood ratio Δl. One can then cut all events producing $\Delta l < t_{\mathrm{cut}}$ for example, to purify the data.

Now, the Neyman-Pearson lemma states that when comparing two simple hypotheses, for a given critical region (cut) and a given significance level, the likelihood ratio is the test statistic that provides the greatest power [152]. In other words, for a simple hypothesis, the likelihood ratio $r_{\mathcal{L}}$ (and any statistic that monotonically increases with $r_{\mathcal{L}}$, such as Δl) gives the greatest chance of finding a true positive for a given chance of finding a false negative. Classifiers based on it would be the most powerful.

The two hypothesis PDFs one would wish to use here are the pulse shapes of the proton recoil and positron annihilation events. These are constructed from binned histograms of PMT hit time residuals. Therefore, if one defines bin i as being centered around time residual t_i, with bin width Δt, the PDF of hypothesis A has a value in bin i of $p_i^A \equiv \int_{t_i-0.5\Delta t}^{t_i+0.5\Delta t} p^A(t)dt$, and similarly for hypothesis B. If one then further defines x_i as the number of measured PMT hits with a time residual in bin

i for an event, the log-likelihood ratio is effectively a dot product, via

$$\begin{aligned} &\Delta l = \vec{p} \cdot \vec{x}, \\ &\vec{x} = \left(x_1,\ x_2,\ ...,\ x_n\right), \\ &\vec{p} = \left(\ln\left(p_1^A - p_1^B\right),\ \ln\left(p_2^A - p_2^B\right),\ ...,\ \ln\left(p_n^A - p_n^B\right)\right). \end{aligned} \tag{7.6}$$

This form is shown here to draw a comparison with the Fisher discriminant later on.

Taking a step back, the hypotheses being tested here are simply the nature of the event, with everything else left undetermined, such as its position in the detector and its energy. However, these quantities were shown to be correlated with the pulse shapes above. Therefore, not only are the measurements not independent, the hypotheses being compared in this classifier are not simple (they are composite), and the Neyman-Pearson lemma thus does not apply. To be more precise, if the event energy were only correlated with the number of measurements (the number of PMT hits here), and this correlation were identical for both hypotheses, the likelihood ratio would be unchanged and the lemma would still apply. However, the shape itself is correlated with energy in non-trivial and differing ways, not to mention the radial position too.

One could produce a PDF for each hypothesis at every energy and radial position (or with as fine a binning as possible), to approach simple hypotheses in each case. One can always reduce a composite hypothesis to a simple hypothesis given enough assumptions. However, this would be a cumbersome process that effectively tries to remove energy and radial dependence without accounting for their inter-correlations. The log-likelihood can of course still be used, but the following method proved more successful.

7.3.2 Fisher Discriminant

An alternative approach to the classification of multi-dimensional data with unknown covariance, originally by Fisher [153], is to find the “best” projection of it to 1 dimension, which allows the easiest separation. This is a type of linear discriminant analysis.

Definition and Motivation

More specifically, consider a training data set with total covariance matrix Σ and total mean $\vec{\mu}$, composed of N multiple underlying data sets, drawn from their own distributions with covariance matrices Σ_i and means $\vec{\mu}_i$ ($i \in \{1, N\}$). Suppose one has n_i data points $\vec{x}_{i,j}$ ($j \in \{1, n_i\}$) in each class, for a total number of data points

$n = \sum_{i=1}^{N} n_i$, the total covariance matrix is

$$\begin{aligned}
\Sigma &= \frac{1}{n}\sum_{i=1}^{N}\sum_{j=1}^{n_i}\left(\vec{x}_{i,j}-\vec{\mu}\right)\left(\vec{x}_{i,j}-\vec{\mu}\right)^T \\
&= \frac{1}{n}\sum_{i=1}^{N}\sum_{j=1}^{n_i}\left(\vec{x}_{i,j}-\vec{\mu}_i+\vec{\mu}_i-\vec{\mu}\right)\left(\vec{x}_{i,j}-\vec{\mu}_i+\vec{\mu}_i-\vec{\mu}\right)^T \\
&= \frac{1}{n}\sum_{i=1}^{N}\sum_{j=1}^{n_i}\left(\vec{x}_{i,j}-\vec{\mu}_i\right)\left(\vec{x}_{i,j}-\vec{\mu}_i\right)^T + \frac{1}{n}\sum_{i=1}^{N} n_i\left(\vec{\mu}_i-\vec{\mu}\right)\left(\vec{\mu}_i-\vec{\mu}\right)^T \\
&= \frac{1}{n}\sum_{i=1}^{N} n_i\Sigma_i + \frac{1}{n}\sum_{i=1}^{N} n_i\left(\vec{\mu}_i-\vec{\mu}\right)\left(\vec{\mu}_i-\vec{\mu}\right)^T \\
\Sigma &= W + B,
\end{aligned} \tag{7.7}$$

where the within-class covariance matrix $W \equiv \sum_{i=1}^{N} \frac{n_i}{n}\Sigma_i$, and the between-class covariance matrix $B \equiv \sum_{i=1}^{N} \frac{n_i}{n}\left(\vec{\mu}_i-\vec{\mu}\right)\left(\vec{\mu}_i-\vec{\mu}\right)^T$ are defined [154]. Now, one can construct the projected ratio of these two

$$R \equiv \frac{\vec{a}^T B \vec{a}}{\vec{a}^T W \vec{a}}, \tag{7.8}$$

which can be maximised by varying the projection vector $\vec{a}$. Note that this quantity is always positive since B and W are covariance matrices (or sums thereof), and so are positive semi-definite. In fact, an additional assumption is made at this stage to get a well defined ratio: W must be positive definite, so that $\vec{a}^T W \vec{a} \neq 0$. This means that $\det(W) \neq 0$ too, and so W is invertible.

Two-Class Solution

Now, in the simpler case of two classes A and B, using $\vec{\mu} = \frac{n_A}{n}\vec{\mu}_A + \frac{n_B}{n}\vec{\mu}_B$ one can easily show

$$\begin{aligned}
B &= \frac{n_A n_B}{n^2}\left(\vec{\mu}_A-\vec{\mu}_B\right)\left(\vec{\mu}_A-\vec{\mu}_B\right)^T, \\
W &= \frac{n_A}{n}\Sigma_A + \frac{n_B}{n}\Sigma_B,
\end{aligned} \tag{7.9}$$

so that the ratio can be written in index notation as

$$R = \frac{n_A n_B}{n^2}\frac{a^i\left(\mu_A^i-\mu_B^i\right)\left(\mu_A^j-\mu_B^j\right)a^j}{a^i W^{ij} a^j}, \tag{7.10}$$

where repeated lower case indices imply summation. Computing $\frac{\partial R}{\partial a^k} = 0$ to find its stationary points, one finds after some algebra

$$\vec{a} = W^{-1}\left(\vec{\mu}_A-\vec{\mu}_B\right)\frac{\vec{a}^T W \vec{a}}{\vec{a}\cdot\left(\vec{\mu}_A-\vec{\mu}_B\right)}, \tag{7.11}$$

which implies $\vec{a} \propto W^{-1}\left(\vec{\mu}_A-\vec{\mu}_B\right)$. Substituting this ansatz back in yields no constraint on the proportionality constant, so the following result is chosen

$$\vec{a} = W^{-1}\left(\vec{\mu}_A-\vec{\mu}_B\right). \tag{7.12}$$

The derivation up to this point can be found in Ref. [152], though it does not "derive" the within-class and between-class covariance matrices (7.7), simply postulating the two-class cases of them (7.9) as useful quantities to work with. It also does not include weighting factors in the definition of W as done here and in Ref. [154] (likewise with B, though it has no impact). This is simply a small difference in the definition of the statistic, which one is free to make. In fact the ratio $r \equiv \frac{n_A}{n_B}$ in the definition of W is the only reasonable free parameter one can insert in the R ratio (7.8): any others simply re-scale R or shift one of the means, and so have no impact. One could simply postulate that W is defined with this parameter, without the previous derivation, and it of course reduces to the definition in Ref. [152] by choosing $r = 1$.

Maximum Proof

To show that this solution is indeed a maximum, the Laplacian of R at the stationary point must be negative[2]. So, calculating $\frac{\partial^2 R}{\partial a^k \partial a^l}$, setting $k = l$ and substituting in the solution for $\vec{a}$ above, one finds

$$\nabla^2 \hat{R} = 2\frac{n_A n_B}{n^2}\left[\frac{(\vec{\mu}_A - \vec{\mu}_B)\cdot(\vec{\mu}_A - \vec{\mu}_B)}{(\vec{\mu}_A - \vec{\mu}_B)^T W^{-1}(\vec{\mu}_A - \vec{\mu}_B)} - \mathrm{Tr}\,(W)\right], \tag{7.13}$$

where $\hat{R}$ denotes R at the stationary point. Rearranging a little, this is

$$\nabla^2 \hat{R} = -2\mathrm{Tr}\,(W)\,\frac{n_A n_B}{n^2}\frac{(\vec{\mu}_A - \vec{\mu}_B)\,M\,(\vec{\mu}_A - \vec{\mu}_B)}{(\vec{\mu}_A - \vec{\mu}_B)^T W^{-1}(\vec{\mu}_A - \vec{\mu}_B)}, \tag{7.14}$$

where

$$M \equiv W^{-1} - \frac{1}{\mathrm{Tr}\,(W)}\mathbf{1}. \tag{7.15}$$

Now, the trace $\mathrm{Tr}\,(W)$ is positive since it is the sum of W's eigenvalues, which are all positive due to W being positive definite. The denominator is also always positive due to W^{-1} being positive definite too, since its eigenvalues are simply the inverses of W's, and thus all positive. Therefore, all that is needed to ensure that $\nabla^2\hat{R} < 0$, is that M be positive definite. To this end, one can find its eigenvalues λ by using the usual formula,

$$\begin{aligned} \det\,(M - \lambda\mathbf{1}) &= 0 \\ \det\left(W^{-1} - \frac{1}{\mathrm{Tr}\,(W)}\mathbf{1} - \lambda\mathbf{1}\right) &= 0 \\ \det\left(W^{-1} - \tilde{\lambda}\mathbf{1}\right) &= 0, \end{aligned} \tag{7.16}$$

where $\tilde{\lambda} \equiv \lambda + 1/\mathrm{Tr}\,(W)$. Notice that for this equation to hold true, $\tilde{\lambda}$ must be an eigenvalue of W^{-1}, which are inverses of W's eigenvalues $\alpha = 1/\tilde{\lambda}$. Thus, one has

$$\lambda = \frac{1}{\alpha} - \frac{1}{\mathrm{Tr}\,(W)} = \frac{\mathrm{Tr}\,(W) - \alpha}{\alpha\mathrm{Tr}\,(W)}, \tag{7.17}$$

---

[2] This part of the proof is not present in Ref. [152], nor could the author find it anywhere else, and so is entirely the author's own work. Though it has likely been worked out somewhere else before. (well done for reading this deep)

where for a particular eigenvalue λ_i $(i \in \{1, n\})$, noting that the trace of a matrix is the sum of its eigenvalues,

$$\lambda_i = \frac{\sum_{j=1}^{n} \alpha_j - \alpha_i}{\alpha_i \sum_{j=1}^{n} \alpha_j}, \tag{7.18}$$

which is always positive since all the eigenvalues are positive, and $n \geq 2$. Therefore M is positive definite, $\nabla^2 \hat{R} < 0$, and the solution found above is indeed always a maximum.

Using the Fisher Discriminant as a Classifier

If one treats with data in the form of multi-dimensional data points $\vec{x}$ (at least 2D), $\vec{a}$ can be computed – or "trained" – on some test data (MC, for example) where the classification is known, using the solution above (7.12). The test statistic can then be computed on real data one wishes to classify, on a point by point basis, via the statistic

$$\mathcal{F} \equiv \vec{a} \cdot \vec{x}. \tag{7.19}$$

One can then sort the data by choosing a threshold $\mathcal{F}(\vec{x}) = \mathcal{F}_{\text{cut}}$. The most obvious threshold is half-way between the projected means

$$t_{\text{cut, Fisher}} = \frac{1}{2} \vec{a} \cdot (\vec{\mu}_A + \vec{\mu}_B), \tag{7.20}$$

but any threshold can be chosen, depending on one's wishes. Notice how similar (7.19) is to the log-likelihood formula (7.6), where $\vec{a}$ has taken on the role of the binned log-likelihood ratio vector $\vec{p}$, while $\vec{x}$ contains the data. This raises the possibility of recasting the one dimensional binned time residual data as a high dimensional vector, by treating each bin as a dimension and its content as the vector component in that dimension. This renders pulse shape discrimination tractable to the Fisher discriminant method.

Link with the Likelihood Ratio

Another interesting link between the Fisher discriminant and the likelihood ratio is shown in Ref. [152], in a limiting case. If one assumes that the PDFs $p^A(x)$ and $p^B(x)$ are both n-dimensional Gaussian distributions for data points $\vec{x}$, with identical covariance matrices Σ

$$p^K(\vec{x}) = \frac{1}{(2\pi)^{n/2} \det(\Sigma)^{1/2}} \exp\left[-\frac{1}{2} (\vec{x} - \vec{\mu}_K)^T \Sigma^{-1} (\vec{x} - \vec{\mu}_K)\right], \tag{7.21}$$

where $K = \{A, B\}$, the log-likelihood ratio becomes

$$\Delta l = (\vec{\mu}_A - \vec{\mu}_B)^T \Sigma^{-1} \vec{x} - \frac{1}{2} \vec{\mu}_A^T \Sigma^{-1} \vec{\mu}_A + \frac{1}{2} \vec{\mu}_B^T \Sigma^{-1} \vec{\mu}_B. \tag{7.22}$$

The two last terms are independent of $\vec{x}$, and so can be ignored as simply shifting the threshold by a constant C. The first term turns out to be exactly Fisher's discriminant (7.19) when the covariance matrices are identical (strictly speaking, proportional to the one defined in (7.12) by a factor of two). So in this simplified case, one finds

$$\Delta l \propto \mathcal{F} + C, \tag{7.23}$$

showing that the Fisher discriminant can also be an "optimal" classifier in certain cases. This equivalence provides no guidance on what an optimal value for the weighting parameter $r = n_A/n_B$ might be however, since it drops out of the equation from the simplification of $\Sigma = \Sigma_A = \Sigma_B$.

7.4 Classifier Training

The word "training" here is used in a loose manner, where it involves selecting appropriate simulation data and formatting it in the best way to be applied to the formulae derived above to compute $\vec{a}$. Note that for now $r = 1$ is assumed, with further discussion of this saved for later.

7.4.1 Reconstructed Distributions

First of all, the same cuts are used on simulated MC data here as are used in the whole present analysis, as shown in table 6.2. An additional maximum prompt energy cut at 3.5 MeV is also applied, to focus only on the energy range where proton recoils occur. Therefore, the energy range looked at throughout this chapter is 0.9-3.5 MeV.

Secondly, there is a risk of the classifier leveraging the different energy spectra of reactor IBDs and (α, n) to separate them, since energy is correlated with the pulse shape. Likewise could be true of the radial position. This is problematic since the energy distribution cannot be treated as known a-priori in an oscillation analysis. Meanwhile the radial distribution of (α, n) events has not been ascertained in detail, and likely changes somewhat over time. The training samples from both should therefore match as closely as possible: have the same energy and radial position distributions. By default, events of any kind are simulated uniformly throughout the full 6 m detector volume, and after reconstruction and cuts look identically distributed.

For the energy distributions, the issue is more complex, since the measured event energy depends on many details of the MC process. The event generators were modified to produce slightly more uniform data, but an extra post-simulation step was still required: all events were grouped into 20 evenly spaced energy bins in the 0.9-3.5 MeV range, and the bin with the least events identified. Events from all other bins were then thrown away at random, until all the bins had the same number of events, producing roughly uniform, and thus similar, distributions. These are shown in figure 7.4.

7.4.2 Constructing the Data Vector

As introduced previously, a vector $\vec{x}$ needs to be constructed for each event, and this can be made by treating each bin of the time residual distribution t_{res} as a dimension, as described for equation (7.6). First however, the t_{res} range and bin width must be decided. To this end, multiple ranges were tested, with a bin width

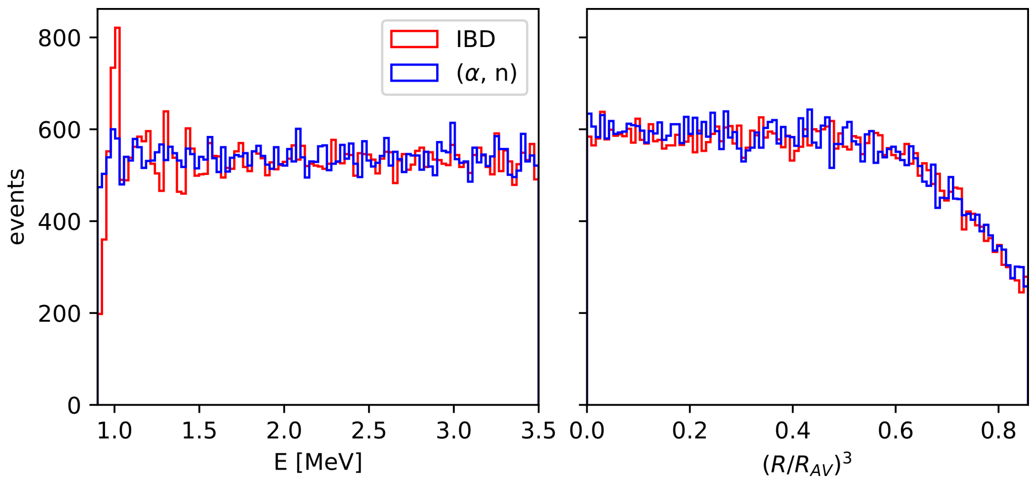

Figure 7.4: Reconstructed energy and radial position distributions of IBD and $(\alpha$, n) prompt events used for classifier training, after cuts and energy spectrum flattening.

fixed at 1 ns, computing $\vec{a}$ in each case, which was then used to classify another set of training data. No change was found by extending beyond $-15 < t_{\text{res}}(\text{ns}) < 150$, so this range was selected. A smaller bin width than 1 ns was not tested, since it is around the limit of the PMTs' resolution. Larger bin widths were tested in the same way, and the classification was extremely robust to these changes: no impact was found up to an 11 ns bin width, and only a minor decrease in performance at 55 ns.

At this point, the Fisher discriminant classifier already out-performed a standard log-likelihood classifier, as shown in figure 7.5. It in fact performed similarly to an improved version of the log-likelihood classifier, where four PDFs were used for each event type, instead of one: split into two energy and two radial position ranges. Its performance can now be enhanced by taking advantage of its use of correlations. The reconstructed radial position R of an event was added as an extra component at the end of $\vec{x}$, yielding the improved classification seen in figure 7.5. Adding the reconstructed energy E instead provided only a very minor improvement, while adding E to the end of the vector already containing R showed no improvement whatsoever. This is because the information about the event's energy is already contained in the magnitude of $\vec{x}$ (excluding the R component), since its magnitude is the number of PMT hits, which was previously shown to be roughly proportional to the event energy. Including only R thus also incorporates any correlation between R and E, explaining why the small improvement of adding E vanishes when R is already there.

Each event thus has a 166-component vector $\vec{x}$, where the last component is its reconstructed radial position, and the rest are residual hit times in the range $-15 < t_{\text{res}}(\text{ns}) < 150$. With this definition, training data was used to compute W and then $\vec{a}$.

7.5 Results

The now "trained" classifier is applied to various simulated and then real datasets, to measure its performance and any biases it might introduce. First its behaviour

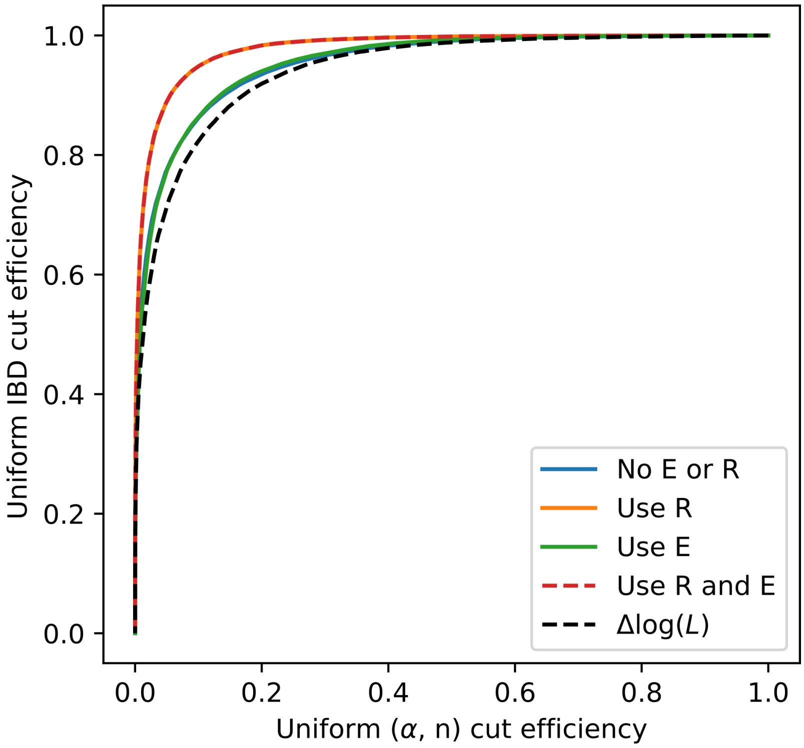

Figure 7.5: Receiver operating characteristic (ROC) curve of the (α, n)/IBD Fisher discriminant classifier, including different information, and applied to simulated training data. Standard and improved log-likelihood classifier showed for comparison.

with uniform training data is further characterised, before more realistic (α, n) and IBD spectra are considered. It is then applied to the real data tagged in chapter 6.

7.5.1 Uniform MC Data

A new dataset of IBD and (α, n) events were simulated with matching energy and radial position spectra in exactly the same way as the training data. These were then classified on an event-by-event basis and the classification $\mathcal{F}$ is shown in figure 7.6 across the prompt energy and radial position ranges. There is correlation between $\mathcal{F}$ and E, as one might expect, since the norm of $\vec{x}$ (sans the last component) is roughly proportional to E. This behaviour is further studied in section 7.6 below. The classifier performs fairly consistently across R ranges, aside from a minor increase in $\mathcal{F}$ at low values.

A cut value was chosen to maximise the signal to background significance

$$\frac{\epsilon_S n_S}{\sqrt{\epsilon_S n_S + \epsilon_B n_B}}, \tag{7.24}$$

where S and B signify signal and background respectively, n is the expected sample size of each type, and ϵ its classification efficiency from the classifier. $n_S = n_B$ is once again assumed, which maximises the significance for a cut at

$$\mathcal{F}_{\text{cut}} = -8.81, \tag{7.25}$$

which yields a signal efficiency of $\epsilon_S = (93.8 \pm 0.1)\,\%$ and a background efficiency of $\epsilon_B = (8.7 \pm 0.1)\,\%$. The uncertainties here are statistical, and estimated by propagating a Binomial standard deviation from the number of selected (not cut) events out of the total: 53280 IBD and 53880 (α, n) events. An interesting point to note is that this is very close to the midway point between the average classifications of the IBDs and (α, n) of -8.85.

A high degree of separation is clearly achieved on average. These classification efficiencies are also shown as a function of E and R^3 in figure 7.6, where the statistical uncertainty in each bin never exceeds around 0.7%. The classifier's performance is relatively constant, aside from an uptick in ϵ_B at low R^3, and a slight decrease in performance at lower energies. For example, below 1.5 MeV it is closer to $\epsilon_S = 90\%$ and $\epsilon_B = 16\%$. The exact classification efficiency will thus depend on the energy spectrum of what is being classified.

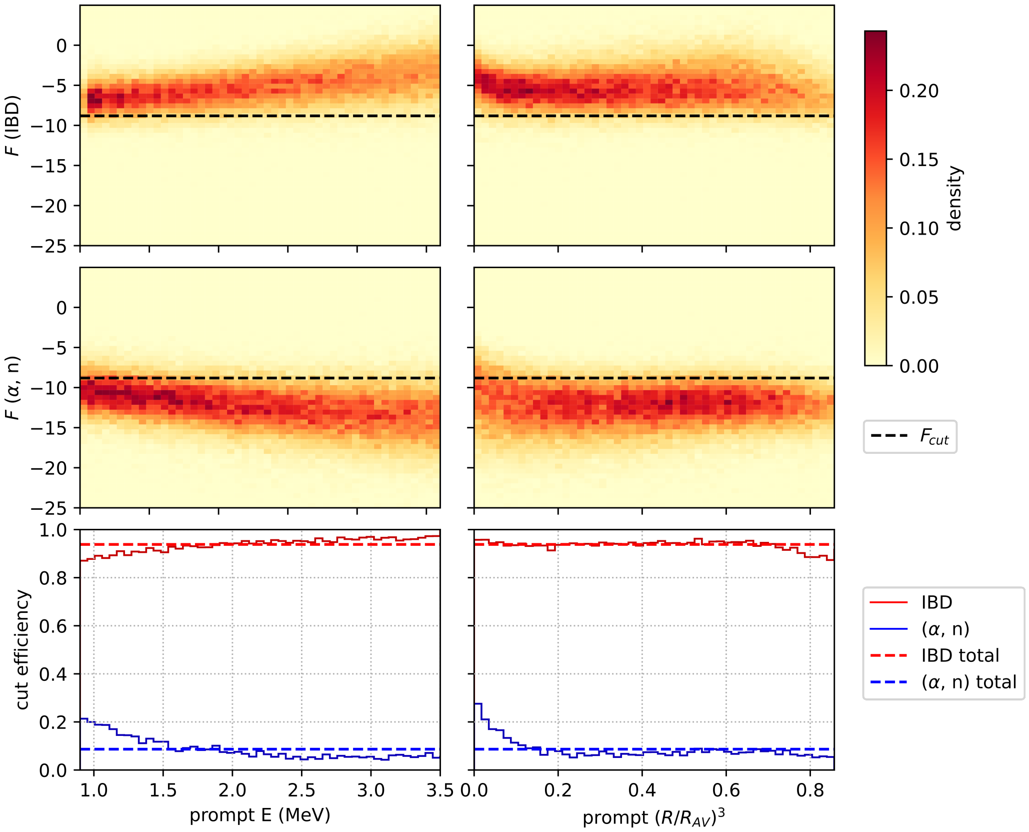

Figure 7.6: Classification F across different energies E and radial positions R, for IBD and (α, n) events issuing from roughly uniform E and R^3 distribution, in the ranges $E \in [0.9, 3.5]$ MeV and $R \in [0, 5.7]$ m. The black dashed line in the top four plots show the cut applied at $\mathcal{F}_{\text{cut}} = -8.81$. The lower two plots show the classification efficiencies of this particular $\mathcal{F}_{\text{cut}}$.

7.5.2 Realistic MC Data

The classifier was then applied to MC simulations of (α, n), geo-ν and reactor IBD events, with the expected energy spectra discussed at length in previous chapters. Just as before, the usual analysis cuts are applied, while only looking at the 0.9 to 3.5 MeV energy range. The antineutrino spectrum is classified both with and without neutrino oscillation applied – using various values of Δm^2_{21} – to see if its effect on the energy spectrum impacts the classifier performance. The results of some of these event-by-event classifications are shown in figure 7.7.

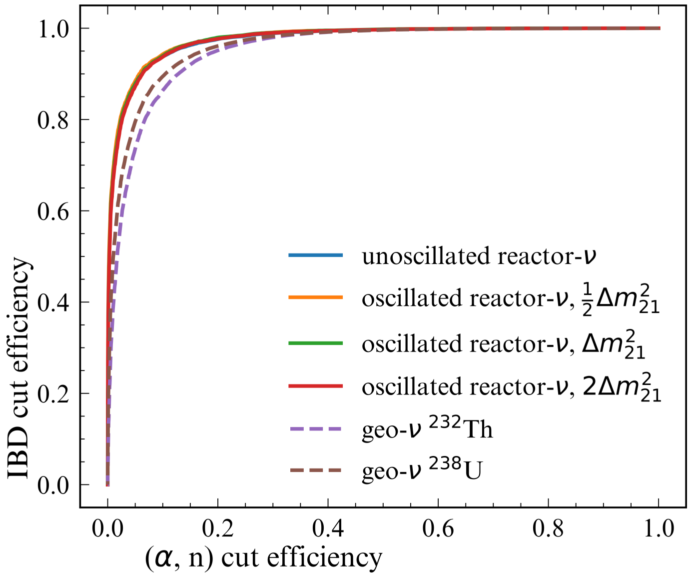

Figure 7.7: ROC curve for MC simulated (α, n) and IBD events issuing from realistic energy distributions. IBD spectra computed with different values of Δm^2_{21} are tested, based on the value from table 5.1.

Neutrino oscillation clearly has no impact on the classifier – the cut efficiency variations are of the same order as the statistical uncertainty – so no complicated oscillation dependence of the signal efficiency need be considered. Meanwhile, the geo-neutrino events have slightly lower efficiencies due to being concentrated around the lowest energies, in which the previous section showed that the IBD classification efficiency is lower. This nevertheless underlines that when the classifier is used in the oscillation analysis, any PDFs should be constructed based on simulations also subject to the classifier, to correctly incorporate its energy dependence.

The exact classifier cut to choose is hard to determine. In principle one would attempt to maximise the signal to background significance again, but this depends on what one considers a signal and a background: reactor neutrinos or geo-neutrinos. Even if one chooses reactor neutrinos, the exact expected number is subject to some uncertainty, particularly if the neutrino oscillation parameters are assumed to be unknown. The expected backgrounds are subject to far more uncertainty still, as discussed in the previous chapter. For the purposes of this analysis, the same cut at $\mathcal{F} = -8.81$ from the previous section is selected, leading to the selection efficiencies in table 7.3. This is not a bad choice anyway, since based on the predictions from table 6.4, the number of reactor IBDs and (α, n) between 0.9 and 3.5 MeV is expected to be roughly equivalent. The accidentals' selection efficiency has not been determined yet, due to time constraints.

7.5.3 Scintillator Timing Test

In order to address the lack of an explicit scintillator timing tuning for protons, a simulation based test was performed. Realistic (α, n) events were simulated and tagged in exactly the same way as the previous section, with the proton timing set

Table 7.3: Selection efficiency (%) of the (α, n)-IBD classifier on signal and backgrounds in the 0.9 to 3.5 MeV range, with statistical uncertainty.

reactor-ν	geo-ν ^{238}U	geo-ν ^{232}Th	(α, n)
94.9 ± 0.3	90.8 ± 0.1	88.1 ± 0.1	11.2 ± 0.3

to either the α or β timings (roughly 4000 events for each) as detailed in section 7.2.2 above. The classification of both simulated datasets are displayed in figure 7.8, where it is clear that it has a completely negligible impact. Specifically the α and β timings lead to classification efficiencies of $(10.9 \pm 0.5)\,\%$ and $(10.2 \pm 0.5)\,\%$, using statistical uncertainties. Based on the discussions in section 7.2.2, the proton timing is expected to be somewhere in-between the two. This clearly demonstrates that it is the event structure of the multiple proton recoils spread out over time that dominates the discrimination power of this classifier. As such, using the α timing as a stand-in for the proton timing is not expected to detract from the classifier's performance.

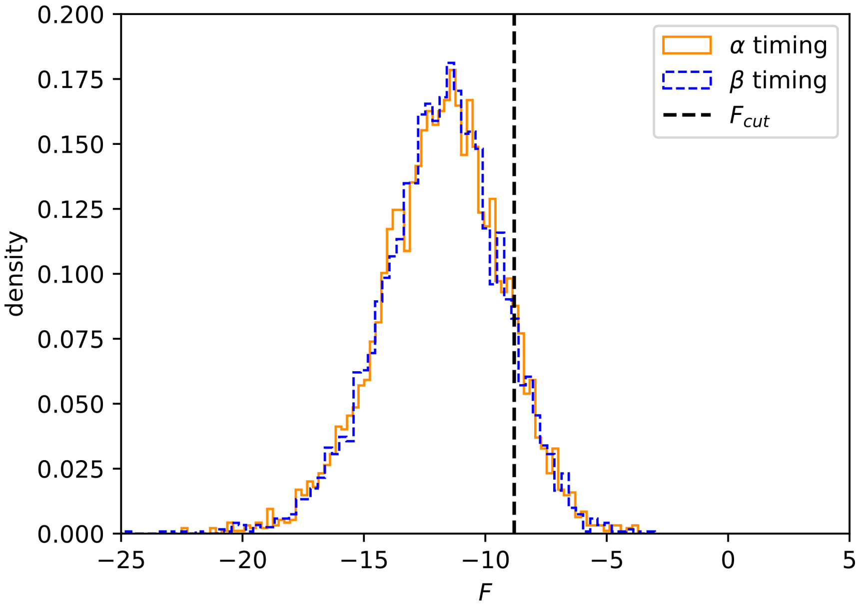

Figure 7.8: Classification $\mathcal{F}$ of simulated (α, n) events issuing from a realistic energy distribution, where the scintillator timing for protons was set either to that of α or β particles. The classifier cut threshold $\mathcal{F}_{\text{cut}} = -8.81$ is also displayed.

7.6 Parameter Tuning

Returning to the previously made assumption of $n_S = n_B$, this is generally the implicit assumption of uses in the literature, unless the values are known and taken from the data. However, in an event-by-event classifier, these quantities cannot be known a priori with certainty. In particular, this classifier is intended to improve the measurement of oscillation parameters, which can strongly depend on the number of measured IBD and (α, n) events. n_S and n_B are thus effectively floating parameters in this analysis.

There is thus no strong mathematical justification for a particular choice of $r \equiv n_S/n_B$, as was further reinforced by the point made in section 7.3.2 above. Instead, different values for this ratio were tested on training data, treating it as a hyper-parameter. As was expected, choosing $n_S > n_B$ grouped the classification of the signal together more strongly, while allowing that of the backgrounds to spread more, and vice-versa. Figure 7.9 bears this out, showing how the standard deviations of the classification of the signal and background vary with r. Indeed, such a choice in r amounts to weighting one within-class covariance matrix more strongly over the other, thus instructing the projection to minimise one's spread at the expense of the other's.

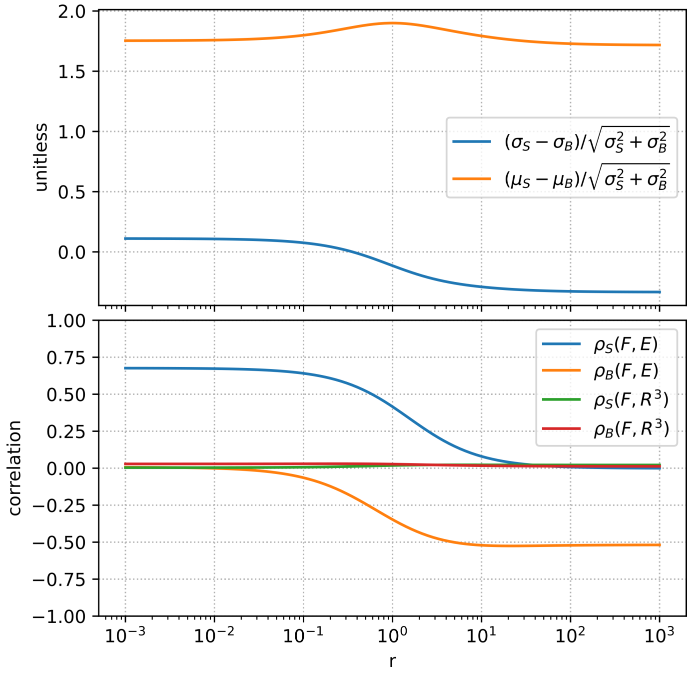

Figure 7.9: Classification of IBD signal and (α, n) background events, issuing from uniform energy spectra, with classifiers using different weightings $r = n_S/n_B$. Subscripts S and B designate signal and background respectively, while $\mathcal{F}$ is the classifier output of an event. For each event type, μ and σ denote the mean and standard deviations in $\mathcal{F}$, while $\rho(\cdot,\cdot)$ denotes the correlation between the two bracketed quantities.

Interestingly, figure 7.9 also shows a strong dependence of the classification-energy correlation $\rho(F, E)$ on the weighting r. Now, non-zero correlation is expected here, since as previously mentioned

$$\mathcal{F} \propto |\vec{x}|, \quad \sum_{i=1}^{n-1} x^i = N_{\text{hit}} \stackrel{\propto}{\sim} E, \tag{7.26}$$

where n is the number of dimensions in $\vec{x}$. However, the chosen projection $\vec{a}$ can be rotated in such a way as to completely suppress this correlation. The catch is that doing so for the signal increases the correlation in the background, and vice-versa. Meanwhile, for equal weighting, the signal and background have equal magnitude but opposite correlations. Such behaviour can be understood thusly: suppose two distributions are plotted in 2-d space, each with their own correlation between the two dimensions. This produces two lines with different slopes. Now, with equal weighting, the best line separating them is a straight line with a slope half-way between those two (the average slope of the two). Rotating this separation line to the x-axis to represent the projection, the two distributions appear the have the same slope magnitude but with opposite signs. Then if one weights one distribution over the other, this amounts to drawing the separation line with a slope closer to that distribution's. Rotating the separation line back to the x-axis again yields a smaller slope for that one at the expense of a greater slope for the other – still with opposite signs.

ROC curves where different values of r were used are shown in figure 7.10. As one might expect, little impact is seen when the IBD and (α, n) events issue from the same uniform E distributions. Even the "real" distributions show only marginally more variation with r, arising from their differing energy spectra. Such behaviour could in principle be exploited to created a further tuned classifier for one's particular needs, if the energy spectra of the data are known a priori. This is not the case here however, and so $r = 1$ is left as-is to avoid any over-tuning of the classifier that may depend on assumptions about the expected energy spectra, and the amount of expected signal and background data.

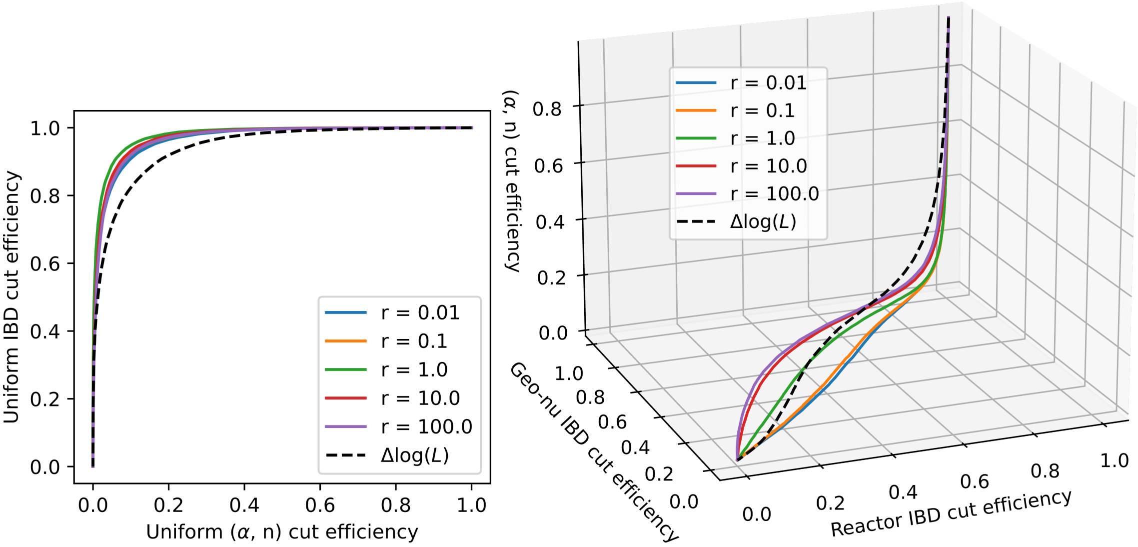

(a) IBD and (α, n) from uniform E distributions.

(b) Reactor IBD, geo-ν IBD and (α, n) from expected E distributions.

Figure 7.10: ROC curves of various signal-background combinations, where various weights r are used in the classifier.

7.7 Real Data

With the classifier's behaviour with MC simulation explored in depth, it can now be applied to data. While characterisation of its performance on known pure data samples is desirable, this has not yet been performed, due both to time constraints and in the case of AmBe data, peculiarities to iron out first. However, the tagged dataset of 57 events presented at the end of the previous chapter is apt for classification. Indeed, this is the goal of the classifier.

The prompt events below 3.5 MeV are classified in figure 7.11a, displaying the mix of IBD-like and (α, n)-like events anticipated in this range. Out of the 48 events in this energy range, 13 are removed by the cut at $\mathcal{F} = -8.81$, which translates to an efficiency of 73%. After neutrino oscillation with the standard PDG value oscillation parameters, about 67% of the predicted (tagged) reactor IBD flux is below 3.5 MeV. Using this combined with the classification efficiencies from table 7.3 and the expected fluxes from table 6.4, the expected classification efficiency of data in the 0.9 to 3.5 MeV range is about 64%. Accidentals were assumed to classify the same as Thorium geo-neutrinos for this, both expected to be low-energy β-like events. Having this classification range from 0% to 100% only changes the result by up to about 1% anyhow. Regardless, given the small sample size and large systematic uncertainties involved, the results appear compatible[3], with a possible small excess around $\mathcal{F} \approx -8$. Furthermore, if the classifier is applied to events above 3.5 MeV, only one out of the nine events is cut – a 3.98 MeV event. This behaviour is also in line with expectation of almost all tagged events above 3.5 MeV being IBD-like, and the classifier performing better for greater energies.

However, the classification of the delayed events tells a mixed story. Aside from the small accidentals contribution, these are all expected to be roughly 2.2 MeV neutron captures, as supported by the distributions shown in 6.5. These events should be IBD-like, being engendered by a single 2.2 MeV γ release, and so almost all tagged delayed events should classify as such. However, figure 7.11b shows that while this trend is roughly followed, there appears to be a systematic shift away from the expected distribution of about -2 $\mathcal{F}$. This discrepancy suggests that simulations are not fully capturing some aspect of the real events, producing subtly different pulse shapes which this classifier teases out.

Finally, while the classification of the prompt events below 3.5 MeV is strongly correlated with their energy as expected ($\rho = 0.53$), it also appears more strongly correlated with the prompt R^3 and ΔR than it should ($\rho = -0.25$ and $\rho = -0.19$ respectively). Not much more can be speculated at this stage, until known pure samples of data with high statistics can be brought to bear, to study the behaviour in detail.

These issues notwithstanding, the cut event pairs (i.e. those with prompt $\mathcal{F} < -8.81$ and prompt $E < 3.5$) are displayed in figure 7.12a, while the resulting prompt energy spectrum is shown in 7.12b, compared to its similarly classified expected distribution. Note that the accidentals PDF is not classified in the latter figure, due to time constraints. These appear to be in good agreement, particularly the prompt energy spectrum. Meanwhile, the predicted (α, n) PR flux has been dramatically reduced.

7.8 Summary

A Fisher discriminant based classifier is constructed, enabling powerful pulse shape discrimination between low-energy prompt IBD events and their main background: (α, n) induced proton recoil events. MC based simulations suggest that 88.8 % of this background can be cut in the 0.9 to 3.5 MeV range, while only sacrificing 5.1 % of reactor IBDs in the same range. The energy and radial position responses of the

---

[3] A note on the data error bars used in this thesis is in appendix A.2

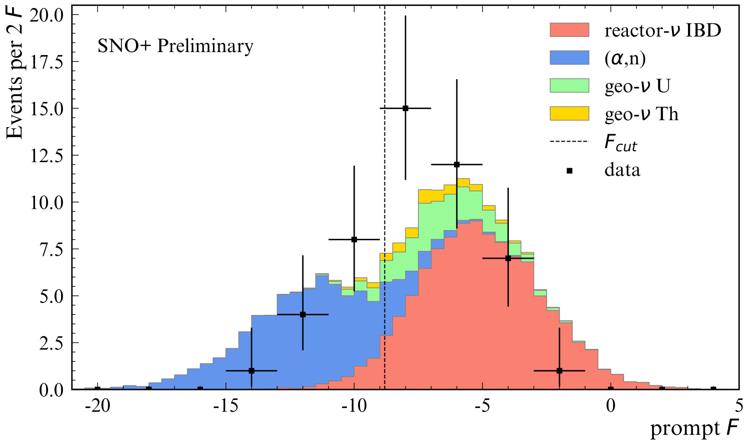

(a) Prompt events below 3.5 MeV, with MC distributions using predicted rates.

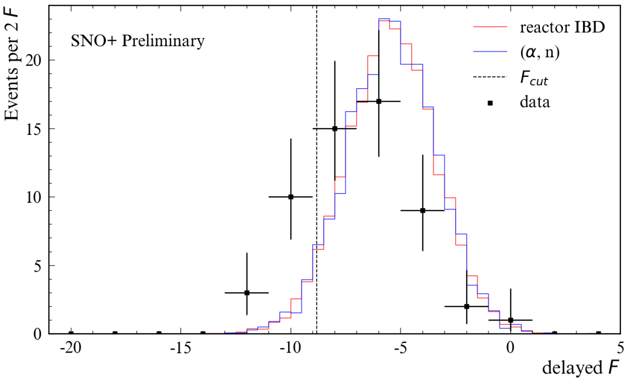

(b) Delayed events, with IBD (α, n) distributions normalised to the number amount of data.

Figure 7.11: Classification $\mathcal{F}$ of tagged prompt and delayed events from data taken between 30/04/2022 and 05/03/2023, compared to their expected distributions from MC simulations.

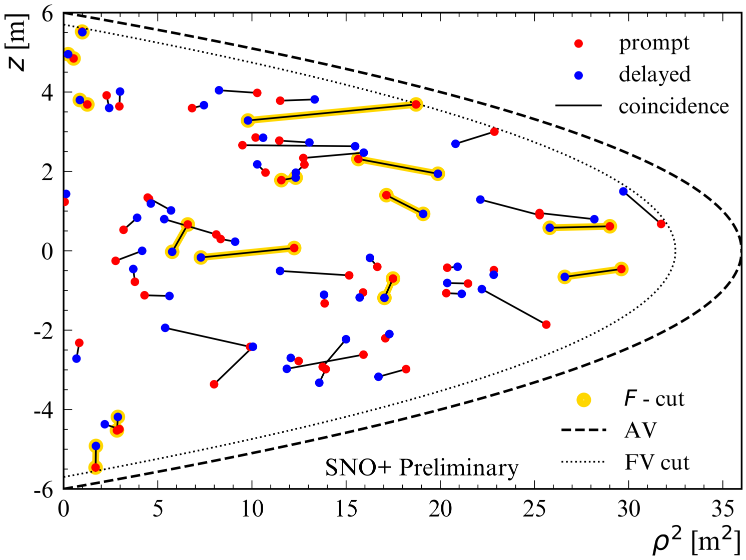

(a) Coincidence pairs displayed according to their reconstructed z and $\rho^2 = x^2 + y^2$ positions in the detector, with those cut by the classifier highlighted in yellow.

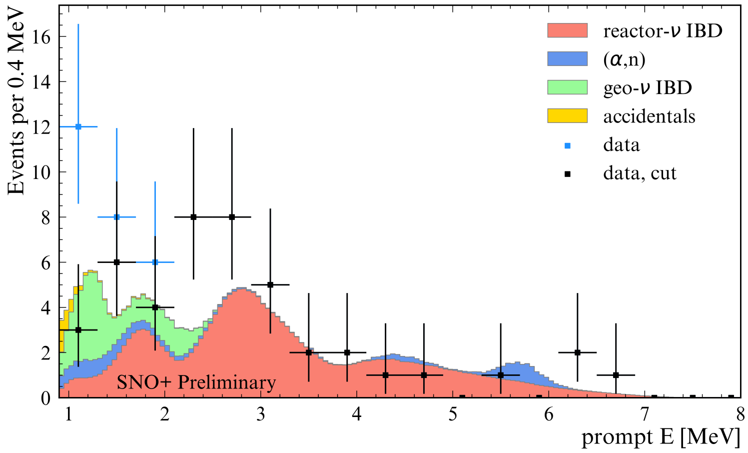

(b) Reconstructed prompt energy, with and without classification, compared to the expected classified spectra (except the accidental PDF which is not classified).

Figure 7.12: Distributions of tagged prompt and delayed events from data taken between 30/04/2022 and 05/03/2023.

classifier were characterised, only inducing minor biases, which can be accounted for by constructing energy spectrum PDFs based on MC simulations subjected to the same classifier.

The characterisation of its performance on real data with large samples has not yet been performed. Plans are in place to use tagged in-situ ^{214}BiPo events to study the classification of IBD-like events, while a future internal deployment of the AmBe source would provide invaluable access to (α, n)-like events. In the meantime, the classifier was applied as-is to the data sample of 57 event pairs under analysis – specifically those with prompt energies below 3.5 MeV. While some potential data-MC mismatch is suggested, the classification results are broadly as expected, and the resulting prompt energy spectrum lines up well with predictions. Definite conclusions are difficult to draw with the current limited sample size. Regardless, the classifier is used in the oscillation analysis at the end of the next chapter, at the very least as a proof of concept.

Systematic uncertainties are hard to assign, though not as critical at this low-statistics stage. A 100% scaling uncertainty is assigned to the classified (α, n) proton recoil PDF, allowing it to scale somewhat independently to the ^{12}C PDF. Otherwise, the geo-neutrinos are already entirely unconstrained, and the accidentals are practically negligible, so no extra uncertainties are assigned to them. The reactor IBD classification only predicts the removal of about 5.1% of them in the low-energy region, which is already covered by highly floating PR and geo-neutrino spectra. Therefore, no classification uncertainty is given to the reactor IBD spectrum either. Only one extra floating parameter is thus to be added to the oscillation analysis when the classifier is employed.

Chapter 8

Oscillation Analysis

"It is a comfort not to be mistaken at all points."
Gandalf – *The Lord of the Rings*, J.R.R. Tolkien

8.1 Introduction

With calculations of the expected reactor antineutrino induced IBD flux and spectrum complete, along with all the most important backgrounds, everything can now be put together to perform a measurement of neutrino oscillation at SNO+. The general principle is to compare data after all the background-reducing cuts to a combination of PDFs of all the expected signal and backgrounds. These PDFs and their corresponding normalisations are then modified for different oscillation parameter values until the best fit is found, via a maximum likelihood analysis. Of course, there are a few more bells and whistles and caveats to this, as will be discussed at length.

Note that only the energy spectrum is studied here, so the PDFs in question are generally one-dimensional. This is due to the expected energy spectra being on the whole well understood, while containing all the important information about Δm^2_{21}. In addition, the amount of data is still quite limited at this point, so that adding other dimensions such as time would not provide much benefit.

The maximum likelihood method is first discussed in detail, followed by an overview of the code package designed to carry out this analysis, and how the various signal and background models are realised therein. Finally, the results are discussed at the end, both with and without using the (α, n)-IBD classifier. In all the methods that follow, it is assumed that the classifier is not used, until the classifier results section. At that point, all the needed modifications to the fitting method are explained.

8.2 Maximum Likelihood

If one wishes to fit data to a PDF that has one or more free parameters (a PDF from a composite hypothesis), one of the most widely used and useful tools is the maximum likelihood (ML) method. It is the tool used in this analysis, so some background information is first given, explaining the final form it will take: the extended binned constrained log-likelihood. How to extract pseudo-confidence limits from this method is then explored.

8.2.1 Building the Extended Binned Constrained Log-Likelihood

The likelihood is defined in essentially the same way here as it was in the context of the likelihood ratio of section 7.3.1, except that it now depends on one or more free parameters $\boldsymbol{\theta}$

$$\mathcal{L}\left(\boldsymbol{\theta}\right) = \prod_{i=1}^{N} p\left(x_i; \boldsymbol{\theta}\right). \tag{8.1}$$

Now, one can fit a model to data by tuning $\vec{\theta}$ to maximise the likelihood

$$\mathcal{L}^{\max} = \mathcal{L}\left(\hat{\boldsymbol{\theta}}\right) = \max_{\boldsymbol{\theta}} \mathcal{L}\left(\boldsymbol{\theta}\right), \tag{8.2}$$

so that $\hat{\boldsymbol{\theta}}$ are the best fit values: the estimators. The ML has some very helpful properties, here taken from Ref. [152]:

- ML estimators are almost always transformation invariant, and unbiased in almost all practical cases.

- ML estimators approach a Gaussian PDF in the large sample limit, a property known as asymptotic normality.

- If an efficient estimator exists, the ML will find it, and ML estimators are always efficient in the large sample limit.

In most practical cases, the log-likelihood $\ln\left(\mathcal{L}(\boldsymbol{\theta})\right)$ is used instead, but all the above properties apply equally to this case since the log-likelihood increases monotonically with the likelihood.

Binned ML

For PDFs arising from non-analytic distributions, such as those produced via Monte-Carlo simulations, one can in practice only use a binned PDF. In this case, the Likelihood is given by the multinomial distribution

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{N!}{\prod_{i=1}^{n_b} N_i!} \prod_{i=1}^{n_b} \left(\frac{\nu_i(\boldsymbol{\theta})}{N}\right)^{N_i}, \tag{8.3}$$

for n_b bins and $N = \sum_{i=1}^{n_b} N_i$, where N_i is the measured number of data points in bin i, with an expected number ν_i. This is a generalisation of the binomial distribution, with the probability of a new data point landing in a bin i given by ν_i/N. The log-likelihood then becomes

$$\ln\left(\mathcal{L}(\boldsymbol{\theta})\right) = \sum_{i=1}^{n_b} N_i \ln\left(\nu_i(\boldsymbol{\theta})\right), \tag{8.4}$$

where any terms not depending on $\boldsymbol{\theta}$ were dropped. In other words, this is effectively a dot product between the data N_i and the log of the PDF $\nu_i(\boldsymbol{\theta})$, as mentioned in section 7.3.1, and it reduces to standard ML in the limit of large data and small bins.

Extended Binned ML

If the number of measurements N also depends on the free parameters $\boldsymbol{\theta}$ one can update the likelihood with the prior Poissonian probability distribution of N,

$$\mathcal{L}(\boldsymbol{\theta}) = \frac{\lambda(\boldsymbol{\theta})^N e^{-\lambda(\boldsymbol{\theta})}}{N!} \frac{N!}{\prod_{i=1}^{n_b} N_i!} \prod_{i=1}^{n_b} \left(\frac{\nu_i(\boldsymbol{\theta})}{\lambda(\boldsymbol{\theta})} \right)^{N_i}, \tag{8.5}$$

where the probability of a data point landing in bin i was changed from ν_i/N to ν_i/λ. $\lambda = \sum_{i=1}^{n_b} \nu_i$ is the total expected number of data points, and so one can simplify this expression to

$$\mathcal{L}(\boldsymbol{\theta}) = e^{-\lambda(\boldsymbol{\theta})} \prod_{i=1}^{n_b} \frac{\nu_i(\boldsymbol{\theta})^{N_i}}{N_i!}, \tag{8.6}$$

which governs the extended binned ML method. The log-likelihood is then simply

$$\ln\left(\mathcal{L}(\boldsymbol{\theta})\right) = -\lambda(\boldsymbol{\theta}) + \sum_{i=1}^{n_b} N_i \ln\left(\nu_i(\boldsymbol{\theta})\right), \tag{8.7}$$

ignoring any terms independent of $\boldsymbol{\theta}$ since only the likelihood's dependence on $\boldsymbol{\theta}$ is relevant.

Constrained Parameters

Consider now the case where the predicted values ν_i and thus λ depend on some other parameters $\boldsymbol{q}$ whose values are known, but not with certainty. Instead, only prior estimates $\boldsymbol{q}_p$ of their values are known, with an associated covariance matrix Σ_q. Therefore, one must vary them to find the maximum likelihood just like $\boldsymbol{\theta}$, but with constraints based on their prior probability distributions. Assuming their uncertainties are Gaussian distributed, the likelihood from (8.6) can once again be updated as

$$\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{q}) = \frac{1}{\sqrt{2\pi |\Sigma_q|}} \exp\left(-\frac{1}{2} \left(\boldsymbol{q} - \boldsymbol{q}_p\right)^T \Sigma_q^{-1} \left(\boldsymbol{q} - \boldsymbol{q}_p\right) \right) e^{-\lambda(\boldsymbol{\theta}, \boldsymbol{q})} \prod_{i=1}^{n_b} \frac{\nu_i(\boldsymbol{\theta}, \boldsymbol{q})^{N_i}}{N_i!}, \tag{8.8}$$

so that the final log-likelihood is

$$\ln\left(\mathcal{L}(\boldsymbol{\theta}, \boldsymbol{q})\right) = -\lambda(\boldsymbol{\theta}, \boldsymbol{q}) + \sum_{i=1}^{n_b} N_i \ln\left(\nu_i(\boldsymbol{\theta}, \boldsymbol{q})\right) - \frac{1}{2} \left(\boldsymbol{q} - \boldsymbol{q}_p\right)^T \Sigma_q^{-1} \left(\boldsymbol{q} - \boldsymbol{q}_p\right), \tag{8.9}$$

and once again constant terms have been discarded. This is the formula that will be used, with one extra detail described below.

Adding PDFs

If the measured PDF is a weighted sum of constituent PDFs, whose weights may themselves depend on $\boldsymbol{\theta}$ and $\boldsymbol{q}$, one can decompose

$$\nu_i(\boldsymbol{\theta}, \boldsymbol{q}) = \sum_{j=1}^{m} \lambda_j(\boldsymbol{\theta}, \boldsymbol{q}) p_j^i(\boldsymbol{\theta}, \boldsymbol{q}), \qquad \lambda(\boldsymbol{\theta}, \boldsymbol{q}) = \sum_{j=1}^{m} \lambda_j(\boldsymbol{\theta}, \boldsymbol{q}). \tag{8.10}$$

p_j^i is the probability of measuring data in bin i according to PDF j, and λ_j is the total expected number of data points from the process described by PDF j. One can substitute these into (8.9), though it is not done in this case. This is because the present analysis first computes the totals in (8.10), before inputting them to the log-likelihood formula (8.9), in two distinct steps. The process is described later, in section 8.3.

8.2.2 Estimator Error and Confidence Intervals

Motivation

Having maximised the likelihood with the estimators $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{q}}$, one can further use the likelihood to assign confidence limits to these estimators. This is done by noting that the ML approximates a Gaussian distribution about the maximum, which can be shown by expanding the log-likelihood about $\hat{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{q}}$. For the rest of this section, the constrained parameters $\boldsymbol{q}$ are lumped in with the unconstrained $\boldsymbol{\theta}$, since they are both maximised and otherwise treated in the same way by the likelihood. The second order Taylor expansion is thus

$$\ln\left(\mathcal{L}(\boldsymbol{\theta})\right) \approx \ln\left(\mathcal{L}(\hat{\boldsymbol{\theta}})\right) + \frac{1}{2}\frac{\partial^2 \ln\left(\mathcal{L}(\boldsymbol{\theta})\right)}{\partial\theta^i\partial\theta^j}\bigg|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}\left(\theta^i - \hat{\theta}^i\right)\left(\theta^j - \hat{\theta}^j\right), \tag{8.11}$$

where the first order term was set to zero, since $\partial \ln\left(\mathcal{L}(\boldsymbol{\theta})\right)/\partial\theta^i = 0$ at the maximum. This means that the likelihood is approximately

$$\mathcal{L}(\boldsymbol{\theta}) \approx \mathcal{L}(\hat{\boldsymbol{\theta}})\exp\left[-\frac{1}{2}\left(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\right)^T \Sigma^{-1}\left(\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}\right)\right], \tag{8.12}$$

where

$$\left(\Sigma^{-1}\right)_{ij} = -\frac{\partial^2 \ln\left(\mathcal{L}(\boldsymbol{\theta})\right)}{\partial\theta^i\partial\theta^j}\bigg|_{\boldsymbol{\theta}=\hat{\boldsymbol{\theta}}}, \tag{8.13}$$

which is also called the Fisher information matrix. Therefore, this Σ takes on the role of the estimators' covariance matrix in the Approximately Gaussian-shaped likelihood. Now, if one defines the log-likelihood ratio similarly to the test statistic in section 7.3.1 as[1]

$$\Delta l\left(\boldsymbol{\theta}\right) \equiv \ln\left(\frac{\mathcal{L}(\boldsymbol{\theta})}{\mathcal{L}(\hat{\boldsymbol{\theta}})}\right), \tag{8.14}$$

this equals zero when $\boldsymbol{\theta} = \hat{\boldsymbol{\theta}}$. Otherwise, if only one parameter deviates away from its ML estimator by N standard deviations, so that $\theta^i = \hat{\theta}^i \pm N\sigma_i$ and $\theta^j = \hat{\theta}^j$ for $i \neq j$, this yields

$$-2\Delta l\left(\hat{\boldsymbol{\theta}} \pm N\sigma_i\right) \approx N^2\sigma_i^2\left(\Sigma^{-1}\right)_{ii}, \tag{8.15}$$

which for uncorrelated estimators $\Sigma = \mathrm{diag}\left(\sigma_1^2, \sigma_2^2, ..., \sigma_m^2\right)$ becomes

$$-2\Delta l\left(\hat{\boldsymbol{\theta}} \pm N\sigma_i\right) \approx N^2. \tag{8.16}$$

Therefore, the log-likelihood difference effectively calculates the N-sigma confidence intervals of uncorrelated estimators.

---

[1] Note that since the components of $\boldsymbol{\theta}$ are not fixed, this is not a simple hypothesis, and so the Neyman-Pearson lemma does not apply. The likelihood ratio is still a very useful quantity however.

2-D Contours

First, the situation with correlated estimators is a little more complex. For example, in the case of two parameters with correlation ρ, equation (8.16) becomes instead

$$-2\Delta l\left(\hat{\boldsymbol{\theta}} \pm N\sigma_i\sqrt{1-\rho^2}\right) \approx N^2, \tag{8.17}$$

so that measured uncertainties are smaller than the true ones by a factor of $\sqrt{1-\rho^2}$. See the discussion in Ref. [155] for more information on the difference between these "inner" and true uncertainties. Competent numerical fitters should be able to find this true uncertainty, by finding the maximum and minimum θ_i that cross the N^2 thresholds while varying all other parameters. Of course, the Fisher information matrix can be computed directly from the log-likelihood space, though it may lose information on asymmetric confidence intervals.

Secondly, if higher dimensional confidence limits are desired, such as two-dimensional confidence contours, the integrated area inside the contour leads to factors different from N^2. These can be calculated by noting from (8.12) that $-2\Delta l\left(\boldsymbol{\theta}\right)$ takes on the form of a χ^2 calculation. In fact, Wilk's theorem states that in the large sample limit, this approximation approaches exactitude [156]. As such, the N-σ limits can be computed by using the quantile of the χ^2 distribution $F^{-1}_{\chi^2}\left(\alpha; n\right)$, which is the inverse of the χ^2 CDF for n degrees of freedom. In the 2-D case, the contours are computed via [155]

$$-2\Delta l\left(\boldsymbol{\theta}_{N-\sigma}\right) = F^{-1}_{\chi^2}\left(2\Phi(N) - 1; 2\right), \tag{8.18}$$

where $\Phi(x) = \int_{-\infty}^{x} \exp(-x'^2/2)dx'/\sqrt{2\pi}$ is the Gaussian distribution's CDF, so that $2\Phi(1) - 1 \approx 0.68$, and $-2\Delta l\left(\boldsymbol{\theta}_{1-\sigma}\right) = 2.30$ for example. This formulation also naturally accounts for correlation between the two parameters in question. Of course, if one takes a slice of the log-likelihood space, so that one parameter is held constant, this reduces to one degree of freedom ($n = 1$), which reproduces the previous N^2 thresholds (and the limitations that come with it).

Caveats

A useful review of different confidence interval constructions is presented in [157], where frequentist-based intervals, such as those described above, are shown to lead to counter-intuitive behaviour at low statistics. Indeed these are better defined as confidence levels (CL), which are statements about expected experimental results, as opposed to the Bayesian-derived confidence intervals (CI) on model parameters. The latter should ideally be used, taking into account the prior probability distribution of possible models. By dealing with frequentist CL to make statements on model parameters, one can end up effectively choosing a prior without clearly stating it.

Furthermore, the CL one finds depends on how they are constructed: the ordering principle, in the words of [157]. The likelihood ratio defines one such ordering principle, proposed by Feldman and Cousins [158] and widely used in the particle physics community. It holds many advantages such as those discussed above, as well as the ability to combine measurements by simply adding log-likelihoods together, but its derived CL can of course suffer from pathologies in the low sample limit. The way nuisance parameters $\boldsymbol{q}$ are dealt with can also muddy the waters, since they may be correlated with the parameters of interest $\boldsymbol{\theta}$. The standard approach is profiling over these nuisance parameters: maximising the likelihood at each value of $\boldsymbol{\theta}$ by varying $\boldsymbol{q}$. This somewhat breaks the frequentist construction if $\boldsymbol{q}$ and $\boldsymbol{\theta}$ are correlated, and introduces some pseudo-Bayesian arguments if prior probability distributions for $\boldsymbol{q}$ are used [157], as they are in the Gaussian constraints defined previously.

Nevertheless, the relative ease of use of the likelihood ratio in otherwise intractable or difficult problems means it is used in the present analysis. In particular, Wilk's theorem, and [157] show that most of these problems disappear with increasing sample size. One must nevertheless bear these caveats in mind, and following the latter's advice, any contours derived here will be called likelihood contours for clarity. Arguments over appropriate confidence contours are left to the still active discussions in the literature.

8.3 `antinuFit` Framework

In order to perform this log-likelihood maximisation for various PDFs, with their complicated dependencies on floating parameters, a suitable framework must be used. The RooFit package was initially tested [159], though it could only effectively handle floating PDF normalisations. It could not be used for the more complicated floating energy systematics, or more importantly the reactor antineutrino oscillation. Instead, an independent code package was developed in `C++`, still using RooFit's numerical minimiser `Minuit`, but otherwise built from the ground up. This `antinuFit` framework turned out to run roughly twice as fast as RooFit under the same conditions, and can handle extra complexities as desired. An overview of its structure and functioning is presented next, with some further areas of improvement and potential extra functionalities pointed out. This was developed with the aim of being passed on and used for successive reactor antineutrino spectral analyses at SNO+, or other similar analyses.

8.3.1 Challenges and Structural Overview

The goal is for this code to be as "plug and play" as possible, while still being computationally efficient and not over-complicated. As such, it was made as modular as feasible, as shown in figure 8.1. The main workhorse is the singleton `Fitter` class, which holds the data, and pointers to all the variables and models in play, and performs the log-likelihood calculation and maximisation.

All variables are saved centrally in a singleton `FitVars` class, so that one variable can be shared by multiple models or energy systematics. This accounts for different models that may depend on some of the same parameters, such as reactor antineutrinos and geo-neutrinos with θ_{12}.

There is also a singleton `Esys` class to hold all the energy systematics transformations. One energy systematic is defined as a single set of energy smearing and scaling, linked to appropriate variables in `FitVars`. In this way, the transformation that is performed changes along with changes to the variables in `FitVars`.

Finally, there are a series of `Models`, which take some input information such as PDFs (Root histograms), constant parameters, and links to relevant variables from `FitVars` and energy systematic from `Esys`. These are used to compute the final expected spectrum (PDF with the correct normalisation) of the model, which is then passed to the `Fitter` class. The spectra from each model are added together and used for the log-likelihood calculation.

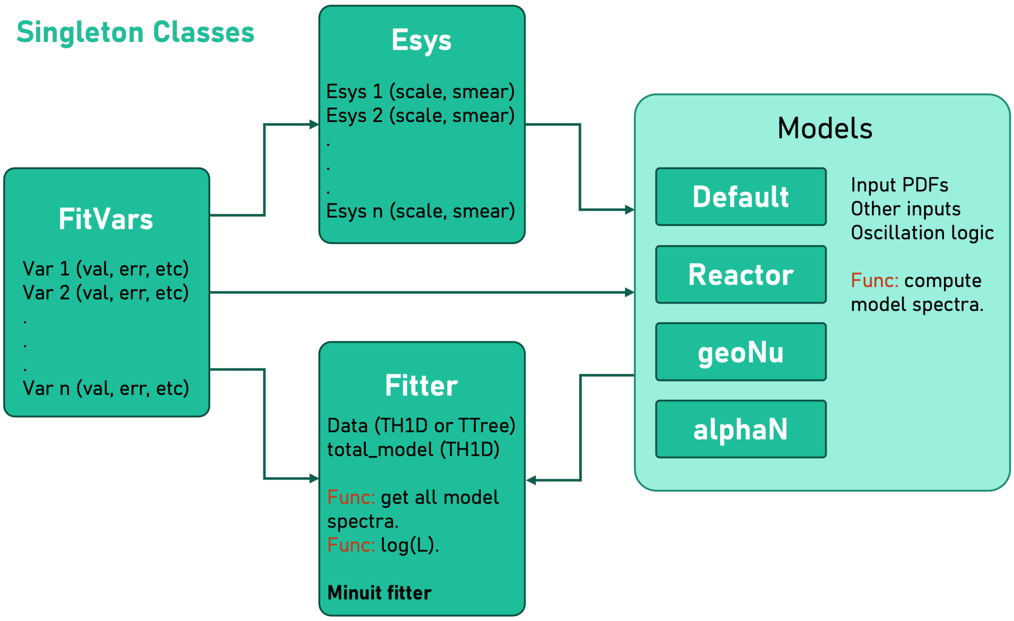

Figure 8.1: Diagrammatic representation of the structure of the `antinuFit` framework.

Putting all this together, the `Minuit` minimiser can run the log-likelihood calculation function – computing all the spectra, and then comparing with the data's event energies to get the result – while varying the variable's values, which are used to produce the spectra. It can therefore minimise the negative of the log-likelihood.

8.3.2 Energy Systematics

Each energy systematic in the `Esys` class is used within a model, and can be shared by as many models as desired. It takes in a PDF in the form of a Root histogram, and performs a transformation to a new PDF as will be described shortly.

First note that all the PDFs discussed here are created using MC simulations subjected to the exact same detector conditions as the data (the same run list), to ensure consistency. Likewise, the same cuts were applied, with one exception: prompt energy cuts, to accommodate systematic uncertainty in reconstructed energy. Recall the systematic energy uncertainties of section 3.7.2, which imply that information from outside the prompt energy cut region may be shifted inside. To implement this appropriately, information about the PDFs beyond these cuts must be saved in so-called sideband regions. As such, the prompt energy limits for PDFs are broader: 0.5 to 9 MeV. The lower limit is beyond the reach of any uncertainty found in this analysis, while all PDFs tend to zero before the upper limit. These extra "padding" regions are shown in an example PDF in figure 8.2. Note that only information that falls inside the 0.9 to 8 MeV range is finally considered after energy systematics have been applied.

The rationale and implementation of systematic uncertainties in the the energy resolution and scaling is laid out below.

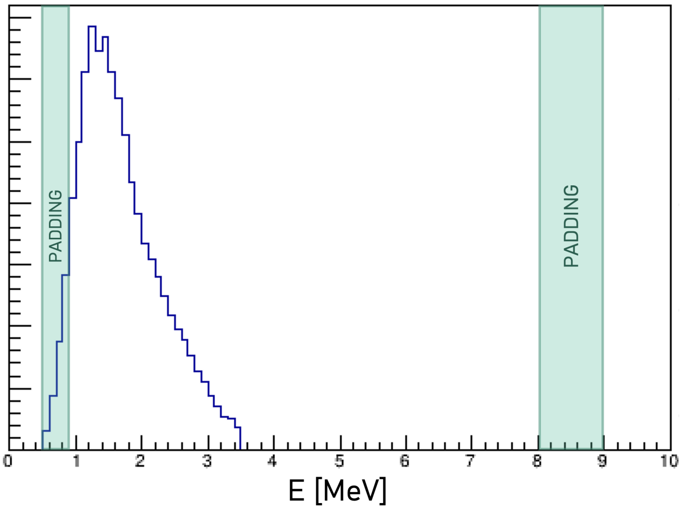

Figure 8.2: Example PDF with extra padding displayed.

Smearing

A systematic uncertainty in the energy resolution was established, which can be simply modelled by applying a Gaussian convolution to the PDFs, with some variable parameter σ. Specifically, this is a transformation from some original distribution $P(E)$ to a smeared distribution $P'(E)$ via

$$P'(E) = \frac{1}{\sigma\sqrt{2\pi}} \int dE' P(E') e^{-\frac{(E-E')^2}{2\sigma^2}}, \tag{8.19}$$

which can be discretised to be used on binned distributions as follows Define $P_i \equiv P(E_i)$, and assume constant bin spacing, so that the bin centers are

$$E_i \equiv i\Delta E + E_{\text{min}}, \tag{8.20}$$

and E_{min} is the bin center for the lowest bin. In this way, the integral can be split up into the bin widths as $E_j - \frac{\Delta E}{2}$ to $E_j + \frac{\Delta E}{2}$, and summed over the bins j. Assuming P_j is constant over this range, and performing a change of variables in the integral, this becomes

$$P'_i = \sum_j M_{ij} P_j, \tag{8.21}$$

where

$$M_{ij} = \Phi\left(\frac{\Delta E}{\sigma}\left[j - i + \frac{1}{2}\right]\right) - \Phi\left(\frac{\Delta E}{\sigma}\left[j - i - \frac{1}{2}\right]\right), \tag{8.22}$$

and $\Phi(x) = \int_{-\infty}^{x} \exp(-x'^2/2)\, dx'/\sqrt{2\pi}$ is the Gaussian distribution's CDF.

However, recall that the smearing uncertainty was defined to be energy dependent, specifically $\sigma(E) = 0.044\sqrt{E}$. As a simplifying assumption, the energy at the centre of the final PDF's bin E_i is used, so that $\sigma \equiv \sigma(E_i)$. Furthermore, for each bin i, computation time is saved by only summing over the j bins whose centres E_j are within 5σ of E_i.

Scaling

Definition Meanwhile, uncertainty in energy scaling is modelled by taking inspiration from Birk's law (7.3), which is parametrised linearly by the scintillator efficiency S, and Birk's constant k_B provides smaller non-linear corrections. As such, the energy scaling is split into two steps: linear scaling and non-linear scaling, defined respectively via

$$P'(E) = P(c \cdot E), \tag{8.23}$$

and

$$P'(E) = P\left(f(E)\right), \quad f(E) = \frac{1 + k_B E}{1 + k'_B E} E. \tag{8.24}$$

Energy scaling is thus achieved by deviations of c away from 1, and of k'_B away from k_B. The non-linear energy scaling form $f(E)$ is originally from [3].

At this point an ordering must be chosen, since $f(cE) \neq cf(E)$; the linear scaling is applied first, so that the overall scaling transformation is governed by $g(E) \equiv f(cE)$, which is

$$g(E) = \frac{1 + k_B c E}{1 + k'_B c E} c E. \tag{8.25}$$

Combining both transformations into one in such a manner also saves on computation time, combining two linear transformations into one.

Now, a scaling uncertainty of 1.8% was found from ^{214}BiPo at 2.2 MeV, which must be split between the linear and non-linear uncertainties. The largest difference in Birk's constants from table 7.2 is 0.004, between the β and proton values. This is thus assigned as the uncertainty in k'_B. Meanwhile, choosing the β's k_B of 0.074, appropriate for the ^{214}Bi decay, an additional uncertainty of 1.1% must be given to c, to produce a total 1.8% uncertainty at 2.2 MeV.

Implementation Once again, this transformation is realised as a matrix transformation of the PDF's bins via

$$P'_i = \sum_j L_{ij} P_j, \tag{8.26}$$

whose computation time can be greatly reduced by only summing over the j bins relevant to each bin i. To do this, one can show that the inverse transformation is

$$g^{-1}(E) = \frac{1}{2k_B c}\left[k'_B E - 1 + \sqrt{\left(1 - k'_B E\right)^2 + 4k_B E}\right], \tag{8.27}$$

which allows one to compute the inverse transformation from the edges of bin P'_i back onto bins P_j. First note that $g(E)$ monotonically increases for any $E > 0$, since $c, k_B, k'_B > 0$, meaning that the lower bin edge of P'_i will always inverse-transform back to a lower point in P_j-space compared to its upper bin edge. Next, writing its bin edges as

$$E_i^{\pm} \equiv E_{\text{min}} + \Delta E \left(i \pm \frac{1}{2}\right), \tag{8.28}$$

one can show that the indices of the bins they land in are

$$j_{\text{min}} = \left\lfloor \frac{g^{-1}\left(E_i^-\right) - E_0^-}{\Delta E} \right\rfloor, \quad j_{\text{max}} = \left\lceil \frac{g^{-1}\left(E_i^+\right) - E_0^+}{\Delta E} \right\rceil, \tag{8.29}$$

where the peculiar brackets are, in order, the floor (round down) and ceiling (round up) functions. If $c = 1$ and $k_B = k'_B$, one can check that this implies $j_{\text{min}} = j_{\text{max}}$, as one might expect.

Now, the L_{ij} factor takes on different values, depending on what fraction of the bin width of P_j is mapped into P'_i, which itself is determined by $j_{\min}$ and $j_{\max}$. If $j_{\min} = j_{\max} = j_{\text{only}}$, both of P'_i's edges are mapped from inside $P_{j_{\text{only}}}$, covering the following fraction of it

$$L_{ij_{\text{only}}} = \frac{1}{\Delta E}\left[g^{-1}\left(E_i^+\right) - g^{-1}\left(E_i^-\right)\right]. \tag{8.30}$$

Otherwise if $j_{\min} \neq j_{\max}$, the lower edge of P_i is mapped into bin $j_{\min}$, meaning the fraction of $P_{j_{\min}}$ mapping into P'_i is instead

$$L_{ij_{\min}} = \frac{1}{\Delta E}\left[E_j^+ - g^{-1}\left(E_i^-\right)\right], \tag{8.31}$$

while the upper edge of P'_i is mapped into bin $j_{\max}$, so that

$$L_{ij_{\max}} = \frac{1}{\Delta E}\left[g^{-1}\left(E_i^+\right) - E_j^-\right]. \tag{8.32}$$

If there are any bins between these limits, they are fully mapped into P'_i, meaning $L_{ij_{\text{between}}} = 1$, where $j_{\min} < j_{\text{between}} < j_{\max}$. One thus only needs to loop over the j bins from $j_{\min}$ to $j_{\max}$ for a particular i, and follow these rules for the values of L_{ij}.

Impact on PDF Normalisation

In principle these energy scalings can change the total measured flux of different signals and backgrounds by moving PDF bins inside or outside the cut region. To disentangle the energy systematics from other normalisation systematics, all the PDFs are scaled to their set normalisations before any systematics are applied. Therefore, the measured normalisation from integrating a resulting spectrum will in general be slightly different from its fit normalisation parameter, even without accounting for neutrino oscillation. Uncertainty propagation to the measured fluxes must also account for this effect in principle. In practice however, they will turn out to have almost no impact.

8.3.3 Models

Each model is its own singleton class, to deal with memory allocation issues related to Root histograms. In the future, these histograms will likely be replaced with a more light-weight option such as vectors. This will allow a parent `Model` class, with a specific sub-class for each tailored model-type, providing an easier way for users to add new models. Nevertheless, models can still be added, and are only taken into account if initialised.

The default `Model` class is the simplest: it takes in one PDF, one energy systematic and one linked normalisation variable; the output spectrum is simply the PDF, with the energy systematics applied and scaled up to the normalisation. Because of this, the `Model` class can in fact hold many models, which are accessed either via an index, or the model's name (given at creation). In this analysis however, it only holds one model: the accidentals spectrum calculation.

All other model classes – the `Reactor`, `alphaN` and `geoNu` models – can only contain one model each, since it is estimated that only one copy of each will ever be necessary. They are also more specialised, taking in more information and performing more tasks in their spectrum calculations, as described below. Note that one energy systematic will be shared by all models, except the `alphaN` model that will turn out to take in an extra independent one too.

Geo-ν IBDs

The two geo-neutrino PDFs – uranium and thorium – are quite straightforwardly simulated using the spectra discussed in the background chapter. The resulting PDFs are displayed in figure 8.3a, after the usual tagging. A total geo-neutrino normalisation is used as an unconstrained floating parameter, while the uranium to thorium ratio is constrained around 3.7 with a 35% uncertainty. The unconstrained total flux also allows for a TNU measurement to be extracted, with no prior assumption.

These two PDFs are taken as inputs to the `geoNu` model, in which the effect of neutrino oscillation is simply applied as an overall scaling factor, using the formula discussed previously. As such, it is the un-oscillated geo-ν flux $N_{\text{geo-}\nu}$ that is given to the model as a variable, not the measured oscillated geo-ν flux. Given that oscillation simply acts a total scaling, and the geo-ν flux is unconstrained anyway, one could in principle use the oscillation flux as the fitter variable instead. The former option was chosen as more natural in this case however.

Finally, this model benefits from a pre-computation function, which allows the user to have the oscillation scaling pre-computed before running the fitting on the rest of the parameters, thus saving computation time.

(α, n)

(α, n) interactions are simulated using a far higher rate than expected, and subjected to the same detector conditions and cuts as data. The resulting spectrum is split into three separate PDFs using simply energy cuts, based on the three peaks caused by the three possible prompt signals discussed in section 6.3.3:

- Proton recoil (PR): below 3.5 MeV.
- ^{12}C scatter (^{12}C): between 3.5 and 5.4 MeV.
- ^{16}O deexcitation (^{16}O): above 5.4 MeV.

These are shown in figure 8.3b, with their expected normalisations. However, the shape and scale of these can change based on a few floating parameters. As discussed in section 6.3.3, the plan here is to roughly follow KamLAND's procedure: A joint normalisation for the first two PDFs (ground state normalisation N_{GS}), with a 30% uncertainty, and a separate normalisation for the third one (excited state normalisation N_{ES}), subject to a 100% uncertainty. The PR PDF is also subjected to an extra, separate linear energy scaling c_P to account for the uncertain proton quenching.

Accidentals

The accidentals PDF is data-driven, as described in the backgrounds chapter. In this case, it is made by binning any events in the dataset that pass all the prompt

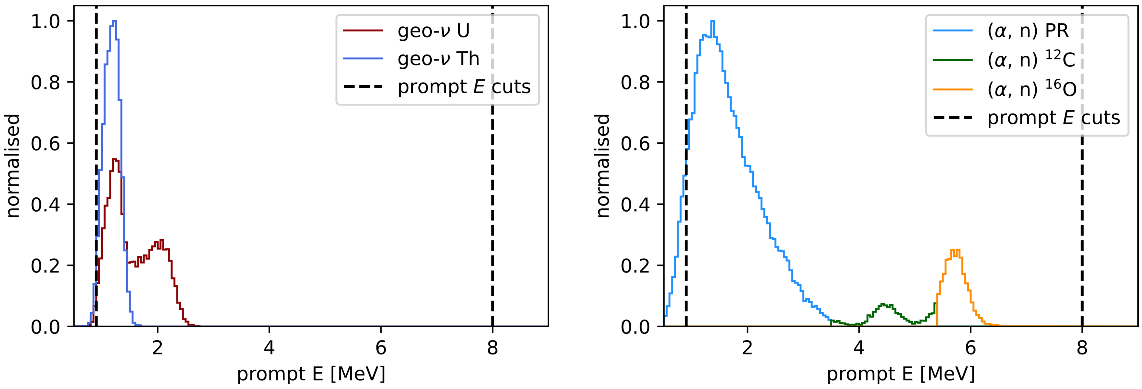

(a) Geo-ν PDFs.

(b) (α, n) PDFs.

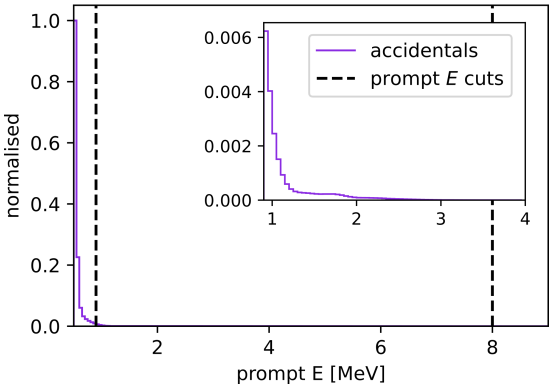

(c) Accidentals PDF, with zoom in at $E \in [0.9, 4]$ MeV.

Figure 8.3: Background PDFs used in the oscillation analysis.

cuts, though with the larger 0.5 to 9 MeV range, and shown in figure 8.3c. It has a floating normalisation N_{acc}, and is the only model with a "trivial" energy systematic (it does nothing), since it was extracted directly from the data. As stated before, this is realised using the default `Model` class.

Reactor-ν IBDs

The `Reactor` model finds itself in the complicated position of needing to produce IBD PDFs in reconstructed prompt event energy (based off the positron energy), while neutrino oscillation impacts them via their true antineutrino energy and the baselines of their originating reactor cores. One solution is to keep all this information saved for a great number of simulated events, and produce the expected PDFs during the oscillation fit for each given set of oscillation parameters. However, given the large number of times such a calculation may need to be performed, a more efficient solution is desirable.

PDFs First, MC simulations were performed, based on the same detector and reactor conditions as the data, but with 30000 times the expected flux. From these the un-oscillated prompt energy spectra, and the true antineutrino energy spectra can be obtained for both PWR and PHWR reactors. These are shown in figure 8.4.

Only two reactor types are modelled here, so only PWR and PHWR spectra are needed. The calculation in section 5.3.4 showed that antineutrinos from reactors

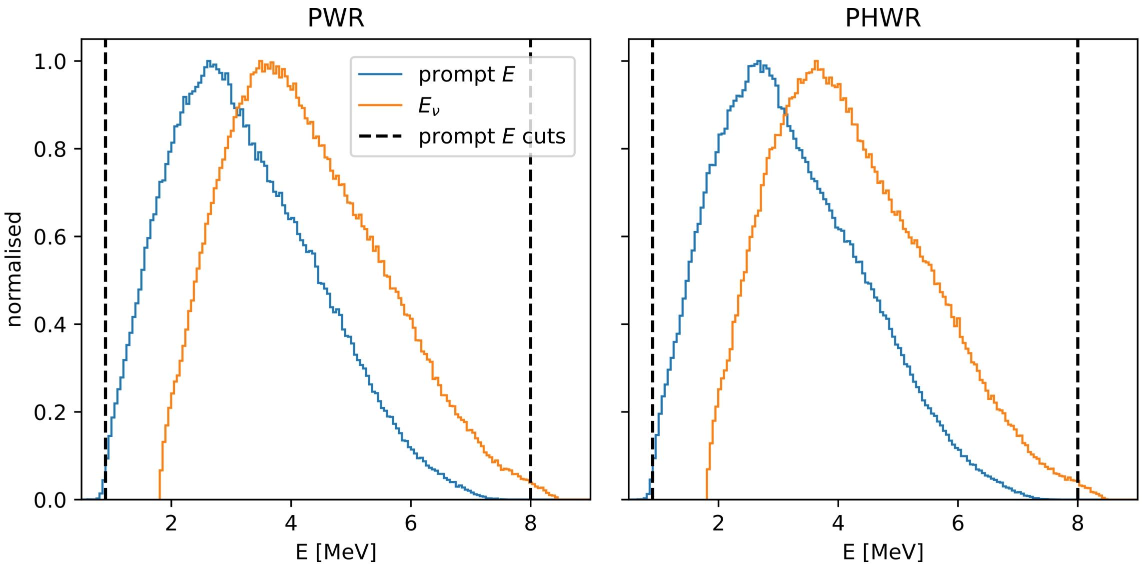

Figure 8.4: Reactor-ν IBD PDFs for PWRs and PHWRs, in both antineutrino energy and prompt energy.

within 1410 km of the SNO+ detector only travel through the Earth's crust, and as such the constant matter density survival probability algorithm will be used. Therefore, the PWR and PHWR spectra in antineutrino energy will be used as PDFs, with the conversion to prompt energy applied later, at the same time as the oscillation. For reactors beyond 1410km, they will simply be modelled as one PWR PDF in prompt energy space. This is because the oscillation becomes fully averaged out a long time before this, and so the same average oscillation scaling as for geo-neutrinos can be used. The vast majority of these reactors are also PWRs or equivalent. Finally, they make up less than 5% of the total expected flux anyway, so any matter effects are negligible for example.

In total, three PDFs are passed to the `Rector` model: PHWR and PWR in antineutrino energy E_ν, and PWR is prompt energy E.

Energy Conversion Notice that the effects of detector and tagging efficiency have been taken into account for all the PDFs above, whether in E_ν or E space. This means that an event with some antineutrino energy E_ν from its governing spectrum has a 100% probability of obtaining some prompt energy E from its own spectrum, and vice-versa.

Further notice that the relationship between the true antineutrino energy E_ν and the reconstructed prompt energy E is independent of the antineutrino's origin. One can thus encode the relationship between the two by constructing a 2-D histogram of these quantities from the MC simulations, as shown in figure 8.5. The x-axis projection of this gives the overall (un-oscillated) IBD spectrum in E-space, while the y-axes projection gives it in E_ν-space. Now, the aforementioned 100% probability argument means that for a particular E_ν, the PDF $p(E|E_\nu)$ of expected E can be computed via the normalisation condition

$$\int_0^\infty p(E|E_\nu)\,dE = 1, \tag{8.33}$$

on the 2-D histogram (except where $p(E|E_\nu) = 0\ \forall E$). Note that this differs from

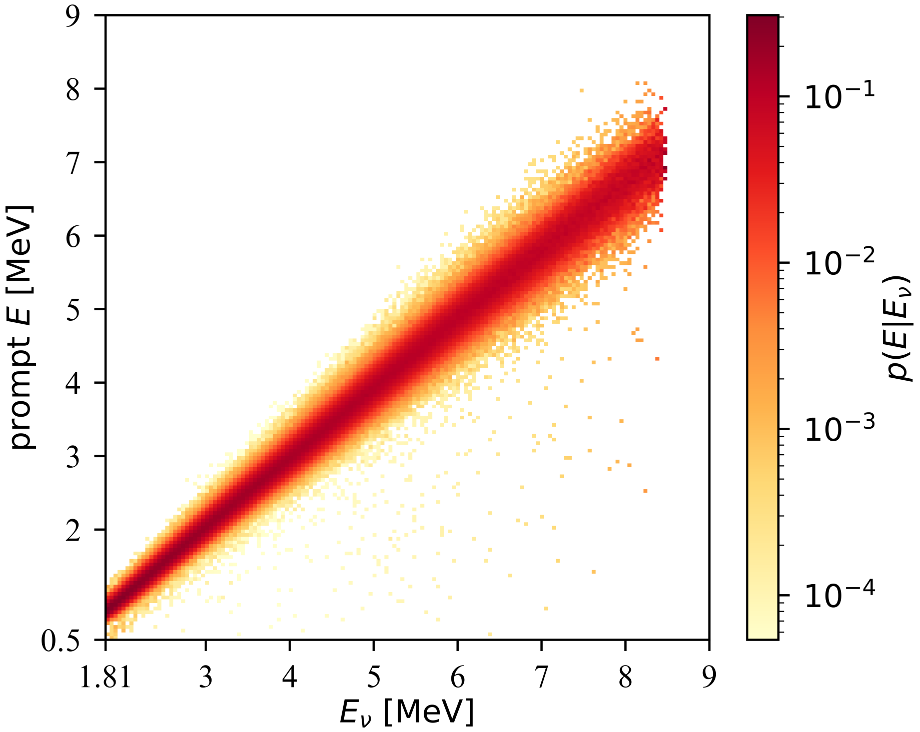

Figure 8.5: 2-D PDF $p\left(E|E_\nu\right)$, which is a PDF of the reconstructed event energy E for every true antineutrino energy E_ν, accounting for detector and tagging efficiencies.

the $p\left(E|E_\nu\right)$ PDF defined earlier in (4.2) by also incorporating the detection and tagging efficiency $\int \eta(E,t)dt$. Likewise, the histogram can provide the the PDF $p\left(E_\nu|E\right)$ of E_ν for a given E by instead enforcing

$$\int_{E_{\text{min}}}^{\infty} p\left(E_\nu|E\right) dE_\nu = 1, \tag{8.34}$$

(again except where $p\left(E_\nu|E\right)\ \forall E_\nu$) In terms of transformation matrices, this is

$$\begin{aligned} M_{ij} &\equiv p\left(E_i|E_j^\nu\right), \quad M_{ij}^{-1} \equiv p\left(E_i^\nu|E_j\right), \\ p\left(E_i|E_j^\nu\right) &= \frac{H_{ij}}{\sum_i H_{ij}}, \quad p\left(E_i^\nu|E_j\right) = \frac{H_{ji}}{\sum_i H_{ji}}, \end{aligned} \tag{8.35}$$

where H_{ij} is the matrix holding the 2-D histogram values, with i indexing the E value (x-axis) and j the E_ν values (y-axis). The second line in these relations are matrix version of the normalisation conditions imposed above. One can check that these indeed obey $\sum_j M_{ij} M_{ij}^{-1} = \sum_j M_{ij}^{-1} M_{ij} = \mathbb{1}$, as they should.

Bringing all this together, the oscillated IBD spectrum from one reactor core n can be computed via

$$N_n(E_i) = N_{\text{IBD}} f_n \sum_j M_{ij} P_{\overline{\nu}_e \to \overline{\nu}_e}\left(E_j^\nu, L_n\right) F_n^\nu(E_j^\nu), \tag{8.36}$$

where f_n is the fraction of un-oscillated flux coming from reactor n, and $F_n^\nu(E_j^\nu)$ is either the PWR or PHWR antineutrino energy spectrum, normalised to unity.

Reactor splitting The Bruce reactor complex is the closest to SNO+, at around 240 km, while the distances to its constituent reactor cores vary by up to 2.7 km. This separation is particularly large because the complex is divided into two sections: Bruce A and Bruce B. As such, this 1.1% baseline variation can impact the oscillation analysis, and so each reactor core at Bruce has oscillation applied separately with its particular baseline. No other complexes withing the first 1000 km approach this reactor core spread. The next closest reactor complex is Pickering, with a maximum spread between its core baselines of around 400 m, or about 0.1% of its average baseline. Such variations were shown to be negligible in the Oscillation chapter. Therefore, all reactor complexes within the earlier mentioned 1410 km limit have oscillation applied to all their constituent cores simultaneously using the average baseline. This is possible since all reactor complexes in this zone are made up of cores that are either all PHWRs or all PWRs. Computation is thus saved, with only negligible impact to the accuracy.

Putting everything together, the total reactor IBD spectrum is computed by the fitting framework as follows

$$\begin{aligned} N(E_i) = N_{\text{reactor-}\nu} \Bigg[\sum_j M_{ij} \Bigg(& F^{\nu}_{\text{PHWR}}(E^{\nu}_j) \sum_n^{\text{Bruce}} f_n P_{\overline{\nu}_e \to \overline{\nu}_e}\left(E^{\nu}_j, L_n\right) \\ & + F^{\nu}_{\text{PHWR}}(E^{\nu}_j) \sum_n^{\text{PHWR}} f_n P_{\overline{\nu}_e \to \overline{\nu}_e}\left(E^{\nu}_j, L_n\right) \\ & + F^{\nu}_{\text{PWR}}(E^{\nu}_j) \sum_n^{\text{PWR}} f_n P_{\overline{\nu}_e \to \overline{\nu}_e}\left(E^{\nu}_j, L_n\right) \Bigg) \\ & + f_{\text{far}} \langle P_{\overline{\nu}_e \to \overline{\nu}_e} \rangle F^{\text{prompt}}_{\text{PWR}}(E_i) \Bigg]. \end{aligned} \tag{8.37}$$

The Bruce sum is over its cores, while the PWR and PHWR sums are over the reactor complexes that are less than 1410 km away from SNO+, aside from Bruce. As such, f_n is the fraction of the un-oscillated flux coming from the selected core or complex, as designated. Meanwhile, f_{far} is the fraction of un-oscillated flux from all nuclear reactors beyond 1410 km, and $F^{\text{prompt}}_{\text{PWR}}(E)$ is the PWR PDF according to prompt energy, instead of the antineutrino energy of the other PDFs $F^{\nu}_{\text{PHWR}}(E^{\nu})$ and $F^{\nu}_{\text{PWR}}(E^{\nu})$.

Note that $N_{\text{reactor-}\nu}$ is the normalisation of the un-oscillated flux, not the oscillated flux. Integrating $N(E)$ would return some fraction of this. $N_{\text{reactor-}\nu}$ is nevertheless the variable that is passed to the fitting process and allowed to vary within constraints, since it is the parameter with a known prior and uncertainty, uncorrelated to the also varying oscillation parameters. Further note that since matter corrections to oscillation parameters described in section 5.5.2 (second step) only depend on E_ν and not on L, these are performed in the energy bin loop (sum over j) to save on computation time. Finally, the vacuum oscillation constants in the pre-computation step (first step) can be pre-computed by the user before fitting to save time, similarly to the `geoNu` model, if the user is controlling the oscillation parameters.

8.3.4 Variables

Each variable saved in the `fitVars` class has a prior, an error, upper and lower bounds, and a current value. In this way the current value can be modified either by the user or during the fitting process, but the prior is stored for any constraint application. Each variable also has two boolians indicating whether it is being held constant (not to be varied by Minuit), and whether it is constrained. A variable being held constant is useful in the context of the user manually testing different combinations of (Δm^2_{21}, θ_{12}) to build a log-likelihood space for example. Meanwhile, if it is constrained, a Gaussian constraint is applied with the prior as its mean and the error as its standard deviation.

This permits one to easily switch between varying (Δm^2_{21}, θ_{12}) manually to produce a log-likelihood space, and then treating them the same as all the other parameters by letting Minuit vary them – whether constrained or not.

Aside from the parameters of interest Δm^2_{21} and θ_{12}, all the other nuisance parameters are passed to the fitter, with the information listed in table 8.1. These are allowed to float with Gaussian constraints, within $\pm 3\sigma$, and only with positive values. Any other numbers are effectively fixed during the fitting, such as the accidentals normalisation N_{acc}, or the reactor IBD fractions f_n.

Table 8.1: Floating (nuisance) parameters in the oscillation analysis, with their estimated values (priors) and systematic uncertainties.

	Description	Symbol	Prior	Uncertainty
Normalisations	geo-ν IBD	$N_{\text{geo-}\nu}$	11.6	unconstrained
	U/Th ratio	$R_{\text{U/Th}}$	3.70	35%
	(α, n) GS	N_{GS}	15.0	30%
	(α, n) ES	N_{ES}	1.50	100%
	reactor-ν IBD	$N_{\text{reactor-}\nu}$	52.2	3%
	accidentals	N_{acc}	0.59	40%
Energy syst.	linear scaling	c	1	1.1%
	linear scaling, PR	c_{P}	1	1.1%
	non-linear scaling	k'_B	0.074 MeV^{-1}	0.004 MeV^{-1}
	smearing	$\sigma/\sqrt{E}$	0 MeV$^{1/2}$	0.042 MeV$^{1/2}$

8.3.5 Minuit

The Minuit minimiser contains many functionalities and different optimisation methods. First for clarity, Minuit was used here, not Minuit2, simply due to software availability. Second, whenever minimisation is performed in the `antinuFit` code, the following Minuit commands [160] are executed in order, with up to 1000 steps in each:

- SIMPLEX: robust multi-dimensional minimiser that does not rely on the minimised function's derivatives, but is not guaranteed to converge on the minimum in a finite amount of time. It is used as a first coarse-scan for the minimum region.
- MIGRAD: most optimal and widely used minimiser in most cases, providing parameter errors, but sensitive to the minimised function's derivatives. Used for final detailed convergence on minimum.

The uncertainties in MIGRAD are calculated along its minimisation path, which may therefore be inaccurate at the minimum if it converged too fast, or if the problem is highly non-linear (asymmetric uncertainties). As such, if uncertainties are requested this is followed by the next two commands in order, again with up 1000 steps each:

- MINOS: asymmetric errors by checking where the minimised function passes some threshold above its minimum. In this case, the threshold is set to the log-likelihood's 0.5.
- HESSE: symmetric, or parabolic, errors from the minimised function's second derivatives. Essentially the same calculation as MIGRAD's errors, but centered on the minimum. This also provides the full covariance matrix of all parameters, and thus their correlations.

Additionally, checks of the minimiser getting stuck or returning nan values are performed between each command above, and the whole process is re-attempted a handful of times. If the process still fails, the user can manually alter some constant parameters, such as Δm^2_{21} or θ_{12}, by a small fraction before attempting this again. In the present analysis, this is exactly what is done in the few such cases, where the parameters are shifted by 10% of the grid scan parameter spacing.

In principle, the HESSE matrix computed numerically above by Minuit could be be calculated manually, using the exact second derivative formula and the minimised parameter values. For example, the uncertainty in Δm^2_{21} could be obtained by calculating the log-likelihood in the exact same way as previously described, but using a function that calculates $\partial^2 P_{\overline{\nu}_e \to \overline{\nu}_e}(E, L)/\partial\left(\Delta m^2_{21}\right)^2$ instead of $P_{\overline{\nu}_e \to \overline{\nu}_e}(E, L)$. The uncertainty in the fitted Δm^2_{21} is then the inverse of this, at the minimum. The long-baseline approximation can even be used for an easier error calculation, as was discussed at the end of the Oscillation chapter. Such a computationally time-saving addition, which would ensure the user knows what exactly is computed, is saved for a potential future improvement.

8.4 Fit Results

The above models were fit to the 57 tagged event energies using the fitting framework and all the parameters described above. Various configurations are employed, primarily in how the long baseline oscillation parameters are treated, but also in whether the (α, n)-IBD classifier is applied. First, results without the use of this classifier are shown, before describing all changes induced by its inclusion, and the resulting fits.

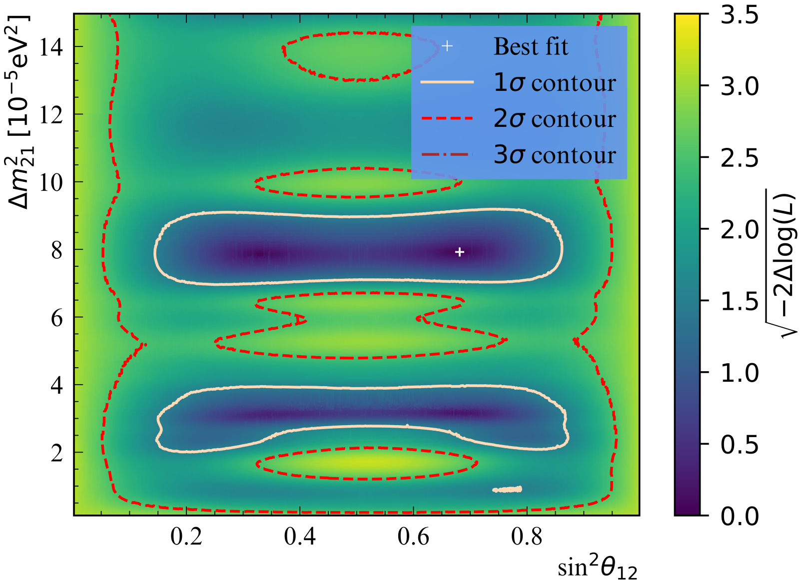

Figure 8.6: Maximised log-likelihood difference phase space, with likelihood contours overlaid.

8.4.1 Without the Classifier

In this regime, the fitting framework is run exactly as previously described, using the floating parameters summarised in table 8.1, with the long baseline oscillation parameters treated in two different ways.

Grid Scan

A grid scan over values of Δm^2_{21} and s^2_{12} was first performed, in ranges of of $[1, 15] \times 10^{-5}$ eV^2 and $[0, 1]$ respectively, split into 500 bins in each dimension. Digressing for a moment, s^2_{12} is used instead of θ_{12}, since it is the former that enters directly in every oscillation formula used herein. One could also use $\sin^2 2\theta_{12}$ due to the approximate symmetry about 45° induced by the long baseline approximation. However, it is in the opinion of the author that at least for an initial study, this approximate symmetry should emerge from the fit rather than being a baked-in assumption. Such parametrisation can of course be more useful for sensitivity and error analyses however.

Nevertheless, at each point in this $(\Delta m^2_{21}, s^2_{12})$ phase space, the fitter is run, so that all other floating parameters are tuned to maximise the log-likelihood. The maximum log-likelihood in the whole grid space is thus the best-fit point, and the the log-likelihood difference across this space is plotted in figure 8.6. The best fit points and the resulting likelihood contours are shown, leading to the following fit values

$$\Delta m^2_{21} = \left(7.92^{+0.82}_{-0.67}\right) \times 10^{-5}\ \text{eV}^2, \quad s^2_{12} = 0.68^{+0.13}_{-0.48}. \tag{8.38}$$

It is clear that these results are far more sensitive to Δm^2_{21} than s^2_{12}, as was expected. This is partly due to the fit θ_{12} depending almost entirely on the overall reactor IBD normalisation, while the number of data points is still limited and the main background fluxes poorly constrained. Additionally, the expected degeneracy about 45° ($s_{12}^2 = 0.5$) is clearly visible, effectively producing the result

$$\sin^2 2\theta_{12} > 0.63. \tag{8.39}$$

Conversely, the fitting of Δm_{21}^2 takes good advantage of the spectral shape and so achieves a better result. Consequently, these two parameters are largely uncorrelated near the best fit, as is evidenced by the 1-σ contour's shape.

The resulting fit prompt energy spectrum is not shown here, but saved for the overall fit below, both for the reasons given therein, but also because the spectra produced are almost identical.

Overall Fits

The method above provides a visually informative log-likelihood space, along with a good estimate for the uncertainties in Δm_{21}^2 and s_{12}^2, by simply noting where they cross the 1-σ contours. Their inter-correlation could also be computed in principle via derivatives of the likelihood space. However, it yields no reliable information on the uncertainties of all other fitted parameters: one can of course extract their best fit values and uncertainties at the best fit point, but this fit held Δm_{21}^2 and s_{12}^2 constant, effectively assuming they were known with absolute precision. This is patently not the case in an oscillation analysis, and so their full correlation with other parameters must be accounted for in order to extract a meaningful uncertainties from all of them.

The versatility of the `antinuFit` framework is very helpful here, allowing Δm_{21}^2 and s_{12}^2 to be instead floated by Minuit, in an overall simultaneous fit of all parameters. First, Δm_{21}^2 and s_{12}^2 are left entirely unconstrained, to get a "SNO+ only" results, essentially assuming uniform priors for them in the considered ranges, if one were to view this in Bayesian terms. Then, gaussian constraints based on the current PDG values were added, obtaining a "PDG + SNO+" result, using all available information. Both results are shown in table 8.2, with uncertainties taken from MINOS and thus potentially asymmetric. The correlations between these fit parameters are then shown in figure 8.7, based on the covariance matrix from MIGRAD, which assumes symmetric uncertainties. Finally, the prompt energy spectra produced by these fits are shown in figure 8.8, where the unconstrained one is almost identical to that produced at the best fit point of the grid scan[2].

The unconstrained oscillation parameter results agree predictably well with those from the grid scan, and agree with the PDG values. Meanwhile the constrained ones are still, largely for Δm_{21}^2, and almost completely for s_{12}^2, dominated by the PDG's prior values. Still, the current measurement of Δm_{21}^2 by SNO+, of which a slightly different fit is to be published very soon, represents the second most precise to date. Meanwhile, the nuisance parameters are largely unchanged between the constrained and unconstrained cases, and do not end up much more constrained after the fit than before. The geo-neutrino flux does, of course, acquire constraints, the most reliable of which are from the constrained fit, so that all available prior information is utilised. Converting the $N_{\text{geo-}\nu} = 19^{+12}_{-11}$ to TNU flux via (6.15) yields

$$\Phi = 70^{+44}_{-40}\ \text{TNU}, \tag{8.40}$$

---

[2] A note on the data error bars used in this thesis is in appendix A.2

Table 8.2: Results of overall fit parameters, compared to their priors, in the cases where the set (Δm^2_{21}, s^2_{12}) is unconstrained and then constrained. *Asymmetric uncertainties reached the maximum allowed range for $N_{\text{geo-}\nu}$: 35.4. The parabolic uncertainty is shown instead.

Parameter	Prior	Unconstrained	Constrained
Δm^2_{21} $[10^{-5}\text{eV}^2]$	7.53 ± 0.18	$8.06^{+0.50}_{-0.44}$	7.59 ± 0.17
s^2_{12}	0.307 ± 0.013	$0.43^{+0.19}_{-0.33}$	0.307 ± 0.013
$N_{\text{geo-}\nu}$	11.8	$23 \pm 13^{*}$	19^{+12}_{-11}
$R_{\text{U/Th}}$	$3.70 \pm 35\%$	$3.78 \pm 33\%$	3.74 ± 34
N_{GS}	$14.91 \pm 30\%$	$15.0 \pm 29\%$	14.8 ± 4.3
N_{ES}	$1.49 \pm 100\%$	$1.00^{+1.24}_{-1.00}$	$0.90^{+1.25}_{-0.90}$
$N_{\text{reactor-}\nu}$	$48.8 \pm 3\%$	$48.90 \pm 2.99\%$	$48.81 \pm 2.98\%$
N_{acc}	$0.59 \pm 40\%$	$0.58 \pm 40\%$	$0.58 \pm 40\%$
c	$1 \pm 1.1\%$	$1.000^{+1.3\%}_{-1.1\%}$	$1.001^{+1.2\%}_{-1.1\%}$
c_{P}	$1 \pm 1.1\%$	$1.000 \pm 1.1\%$	$1.000 \pm 1.1\%$
k'_B $[\text{MeV}^{-1}]$	0.074 ± 0.004	$0.073^{+0.003}_{-0.005}$	$0.073^{+0.003}_{-0.005}$
$\sigma/\sqrt{E}$ $[\text{MeV}^{1/2}]$	0 ± 0.042	$0^{+0.046}_{-0.042}$	$0^{+0.047}_{-0.042}$

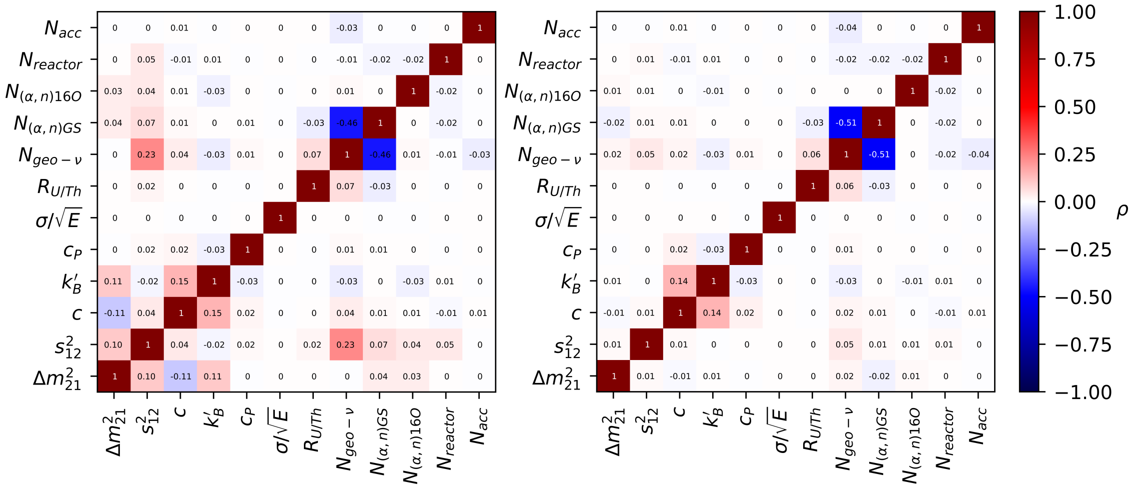

	Δm^2_{21}	s^2_{12}	c	k'_B	c_P	$\sigma/\sqrt{E}$	$R_{U/Th}$	$N_{geo-\nu}$	$N_{(\alpha,n)GS}$	$N_{(\alpha,n)16O}$	$N_{reactor}$	N_{acc}
N_{acc}	0	0	0.01	0	0	0	0	-0.03	0	0	0	1
$N_{reactor}$	0	0.05	-0.01	0.01	0	0	0	-0.01	-0.02	-0.02	1	0
$N_{(\alpha,n)16O}$	0.03	0.04	0.01	-0.03	0	0	0	0.01	0	1	-0.02	0
$N_{(\alpha,n)GS}$	0.04	0.07	0.01	0	0.01	0	-0.03	-0.46	1	0	-0.02	0
$N_{geo-\nu}$	0	0.23	0.04	-0.03	0.01	0	0.07	1	-0.46	0.01	-0.01	-0.03
$R_{U/Th}$	0	0.02	0	0	0	0	1	0.07	-0.03	0	0	0
$\sigma/\sqrt{E}$	0	0	0	0	0	1	0	0	0	0	0	0
c_P	0	0.02	0.02	-0.03	1	0	0	0.01	0.01	0	0	0
k'_B	0.11	-0.02	0.15	1	-0.03	0	0	-0.03	0	-0.03	0.01	0
c	-0.11	0.04	1	0.15	0.02	0	0	0.04	0.01	0.01	-0.01	0.01
s^2_{12}	0.10	1	0.04	-0.02	0.02	0	0.02	0.23	0.07	0.04	0.05	0
Δm^2_{21}	1	0.10	-0.11	0.11	0	0	0	0	0.04	0.03	0	0

	Δm^2_{21}	s^2_{12}	c	k'_B	c_P	$\sigma/\sqrt{E}$	$R_{U/Th}$	$N_{geo-\nu}$	$N_{(\alpha,n)GS}$	$N_{(\alpha,n)16O}$	$N_{reactor}$	N_{acc}
N_{acc}	0	0	0.01	0	0	0	0	-0.04	0	0	0	1
$N_{reactor}$	0	0.01	-0.01	0.01	0	0	0	-0.02	-0.02	-0.02	1	0
$N_{(\alpha,n)16O}$	0.01	0.01	0	-0.01	0	0	0	0	0	1	-0.02	0
$N_{(\alpha,n)GS}$	-0.02	0.01	0.01	0	0	0	-0.03	-0.51	1	0	-0.02	0
$N_{geo-\nu}$	0.02	0.05	0.02	-0.03	0.01	0	0.06	1	-0.51	0	-0.02	-0.04
$R_{U/Th}$	0	0	0	0	0	0	1	0.06	-0.03	0	0	0
$\sigma/\sqrt{E}$	0	0	0	0	0	1	0	0	0	0	0	0
c_P	0	0	0.02	-0.03	1	0	0	0.01	0	0	0	0
k'_B	0.01	0	0.14	1	-0.03	0	0	-0.03	0	-0.01	0.01	0
c	-0.01	0.01	1	0.14	0.02	0	0	0.02	0.01	0	-0.01	0.01
s^2_{12}	0.01	1	0.01	0	0	0	0	0.05	0.01	0.01	0.01	0
Δm^2_{21}	1	0.01	-0.01	0.01	0	0	0	0.02	-0.02	0.01	0	0

Figure 8.7: Correlation matrices of fit parameters in the cases where Δm^2_{21} and θ_{12} are unconstrained (left) and constrained (right).

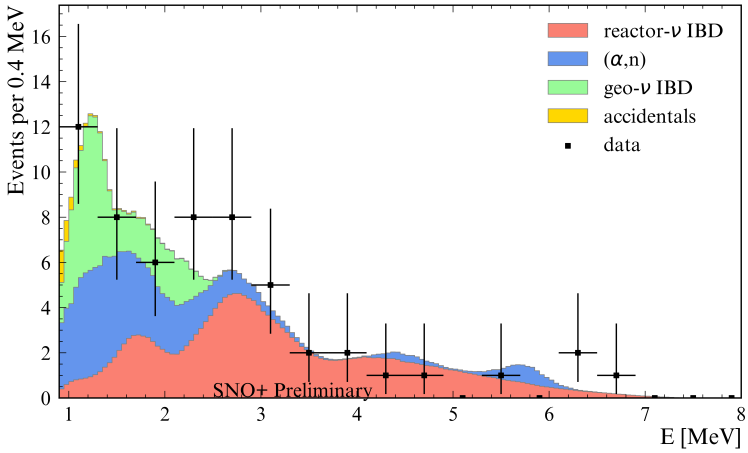

(a) Unconstrained.

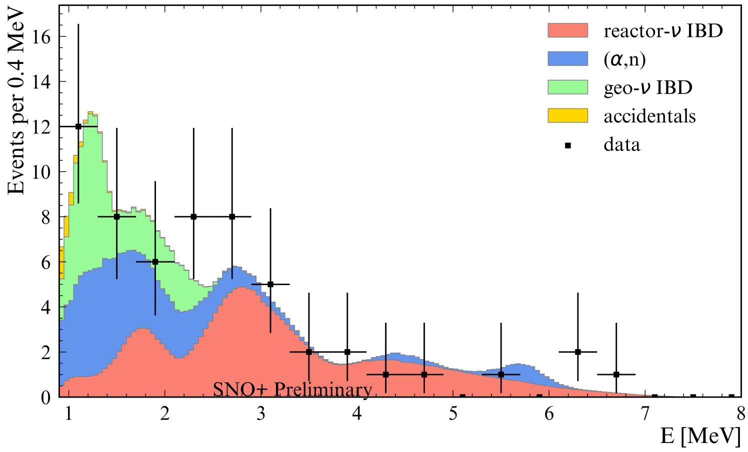

(b) Constrained.

Figure 8.8: Fit prompt energy spectra from oscillation analyses with unconstrained and constrained Δm^2_{21} and s^2_{12}.

which is the third location worldwide that such a measurement has been made. Preliminary measurements of Δm^2_{21} and the TNU flux were presented at the Neutrino2024 conference, based on this work and on that of [3].

The correlation matrices in figure 8.7 show how the correlation between the geo-neutrino flux and s^2_{12} is eliminated by constraining the latter, as one might expect. In fact almost all correlations are reduced to negligible levels, both from the normalisation impact of s^2_{12} and the energy scaling impact of Δm^2_{21}. The one major exception to this is the correlation between the (α, n) ground state and geo-ν normalisations, since they are for the most part on top of each-other. Proton recoil (α, n) events are thus the last major obstacle to more accurate TNU flux measurements at SNO+.

8.4.2 With the Classifier

The (α, n)-IBD classifier appears poised to eliminate the greatest remaining obstacle to these fits, statistics aside.

Fit Results

As such, the same fits are performed once again, while using both data and MC constructed PDFs subjected to the classifier; that is, any tagged prompt events with $E < 3.5$ MeV and $\mathcal{F} < -8.81$ are rejected. The priors of the (α, n) GS and geo-ν normalisations were adjusted according to the classification efficiencies in 7.3, while the efficiency on un-oscillated reactor IBDs in the affected energy region scales its total normalisation by 96.1%. Furthermore, an additional scaling factor S_{PR} is applied to the resulting PR PDF, with a prior of 1 and a Gaussian constrained uncertainty of 100%, as decided at the end of chapter 7.

The grid scan failed to produce reasonably well-fit spectra, and so is only shown in the appendix for the curious reader. It essentially fit to the secondary minima shown around 3×10^{-5} eV^2 in figure A.1, producing visibly worse prompt energy fits. The overall fits however produced more robust results, and are presented in the same way as the previous section: the fit parameters are listed in table 8.3 for both the constrained an unconstrained cases, while their correlations are shown in figure 8.9, and the resulting spectra in 8.10.

In general the results here are very similar to those without the classifier, though interestingly the anti-correlation between the geo-ν flux and s^2_{12} is turned into a strong positive correlation, in the unconstrained case. The strong correlation between the proton recoil flux and the geo-ν still remains after both constraints and the classifier are applied, while the constrained fit of the geo-ν flux yields

$$\Phi = 48 \pm 30 \text{ TNU}, \tag{8.41}$$

accounting for the classifier efficiency in the conversion. This and other uncertainties are marginally larger than previously, indicating that the reduction in data has a larger impact than the reduction in background. In other words, the analysis is currently statistics limited, and reductions in systematic uncertainty yield comparatively little benefit for the time being.

Table 8.3: Parameters from overall fits that used the (α, n)-IBD classifier, compared to their priors, in the cases where the set (Δm^2_{21}, s^2_{12}) is unconstrained and then constrained. *Asymmetric uncertainties reached the maximum allowed range for $N_{\text{geo-}\nu}$: 31.8. The parabolic uncertainty is shown instead.

Parameter	Prior	Unconstrained	Constrained
Δm^2_{21} [10^{-5}eV^2]	7.53 ± 0.18	$7.83^{+0.76}_{-0.61}$	7.56 ± 0.17
s^2_{12}	0.307 ± 0.013	$0.23^{+0.17}_{-0.11}$	0.306 ± 0.013
$N_{\text{geo-}\nu}$	10.6	$9.89^{+11.48}_{-8.03}$	$12.9 \pm 8.2^{*}$
$R_{\text{U/Th}}$	$3.70 \pm 35\%$	$3.88^{+33\%}_{-32\%}\%$	$3.92^{+32\%}_{-31\%}$
N_{GS}	$2.61 \pm 30\%$	$2.56 \pm 30\%$	$2.57 \pm 30\%$
S_{PR}	$1 \pm 100\%$	$1.08^{+99\%}_{-98\%}$	$1.09^{+99\%}_{-98\%}$
N_{ES}	$1.17 \pm 100\%$	$0.80^{+1.04}_{-0.80}$	$0.83^{+1.03}_{-0.83}$
$N_{\text{reactor-}\nu}$	$46.89 \pm 3\%$	$47.00^{+2.99\%}_{-2.98\%}$	$47.11 \pm 2.96\%$
N_{acc}	$0.59 \pm 40\%$	$0.57 \pm 41\%$	$0.57 \pm 40\%$
c	$1 \pm 1.1\%$	$1.000 \pm 1.1\%$	$1.001 \pm 1.1\%$
c_{P}	$1 \pm 1.1\%$	$1.000 \pm 1.1\%$	$1.000 \pm 1.1\%$
k'_B [MeV^{-1}]	0.074 ± 0.004	0.073 ± 0.004	0.073 ± 0.004
$\sigma/\sqrt{E}$ [MeV$^{1/2}$]	0 ± 0.042	$0^{+0.049}_{-0.042}$	$0^{+0.050}_{-0.042}$

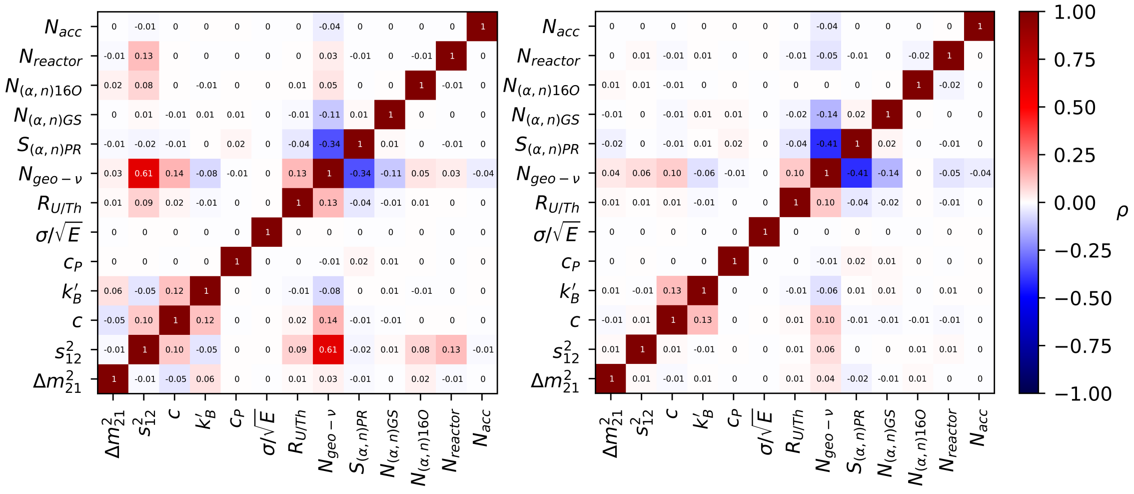

	Δm^2_{21}	s^2_{12}	c	k'_B	c_P	$\sigma/\sqrt{E}$	$R_{U/Th}$	$N_{geo-\nu}$	$S_{(\alpha,n)PR}$	$N_{(\alpha,n)GS}$	$N_{(\alpha,n)16O}$	$N_{reactor}$	N_{acc}
N_{acc}	0	-0.01	0	0	0	0	0	-0.04	0	0	0	0	1
$N_{reactor}$	-0.01	0.13	0	0	0	0	0	0.03	-0.01	0	-0.01	1	0
$N_{(\alpha,n)16O}$	0.02	0.08	0	-0.01	0	0	0.01	0.05	0	0	1	-0.01	0
$N_{(\alpha,n)GS}$	0	0.01	-0.01	0.01	0.01	0	-0.01	-0.11	0.01	1	0	0	0
$S_{(\alpha,n)PR}$	-0.01	-0.02	-0.01	0	0.02	0	-0.04	-0.34	1	0.01	0	-0.01	0
$N_{geo-\nu}$	0.03	0.61	0.14	-0.08	-0.01	0	0.13	1	-0.34	-0.11	0.05	0.03	-0.04
$R_{U/Th}$	0.01	0.09	0.02	-0.01	0	0	1	0.13	-0.04	-0.01	0.01	0	0
$\sigma/\sqrt{E}$	0	0	0	0	0	1	0	0	0	0	0	0	0
c_P	0	0	0	0	1	0	0	-0.01	0.02	0.01	0	0	0
k'_B	0.06	-0.05	0.12	1	0	0	-0.01	-0.08	0	0.01	-0.01	0	0
c	-0.05	0.10	1	0.12	0	0	0.02	0.14	-0.01	-0.01	0	0	0
s^2_{12}	-0.01	1	0.10	-0.05	0	0	0.09	0.61	-0.02	0.01	0.08	0.13	-0.01
Δm^2_{21}	1	-0.01	-0.05	0.06	0	0	0.01	0.03	-0.01	0	0.02	-0.01	0

	Δm^2_{21}	s^2_{12}	c	k'_B	c_P	$\sigma/\sqrt{E}$	$R_{U/Th}$	$N_{geo-\nu}$	$S_{(\alpha,n)PR}$	$N_{(\alpha,n)GS}$	$N_{(\alpha,n)16O}$	$N_{reactor}$	N_{acc}
N_{acc}	0	0	0	0	0	0	0	-0.04	0	0	0	0	1
$N_{reactor}$	0	0.01	-0.01	0.01	0	0	-0.01	-0.05	-0.01	0	-0.02	1	0
$N_{(\alpha,n)16O}$	0.01	0.01	-0.01	0	0	0	0	0	0	0	1	-0.02	0
$N_{(\alpha,n)GS}$	-0.01	0	-0.01	0.01	0.01	0	-0.02	-0.14	0.02	1	0	0	0
$S_{(\alpha,n)PR}$	-0.02	0	-0.01	0.01	0.02	0	-0.04	-0.41	1	0.02	0	-0.01	0
$N_{geo-\nu}$	0.04	0.06	0.10	-0.06	-0.01	0	0.10	1	-0.41	-0.14	0	-0.05	-0.04
$R_{U/Th}$	0.01	0.01	0.01	-0.01	0	0	1	0.10	-0.04	-0.02	0	-0.01	0
$\sigma/\sqrt{E}$	0	0	0	0	0	1	0	0	0	0	0	0	0
c_P	0	0	0	0	1	0	0	-0.01	0.02	0.01	0	0	0
k'_B	0.01	-0.01	0.13	1	0	0	-0.01	-0.06	0.01	0.01	0	0.01	0
c	-0.01	0.01	1	0.13	0	0	0.01	0.10	-0.01	-0.01	-0.01	-0.01	0
s^2_{12}	0.01	1	0.01	-0.01	0	0	0.01	0.06	0	0	0.01	0.01	0
Δm^2_{21}	1	0.01	-0.01	0.01	0	0	0.01	0.04	-0.02	-0.01	0.01	0	0

Figure 8.9: Correlation matrices of fit parameters in the cases where Δm^2_{21} and θ_{12} are unconstrained (left) and constrained (right), and the (α, n)-IBD classifier was used.

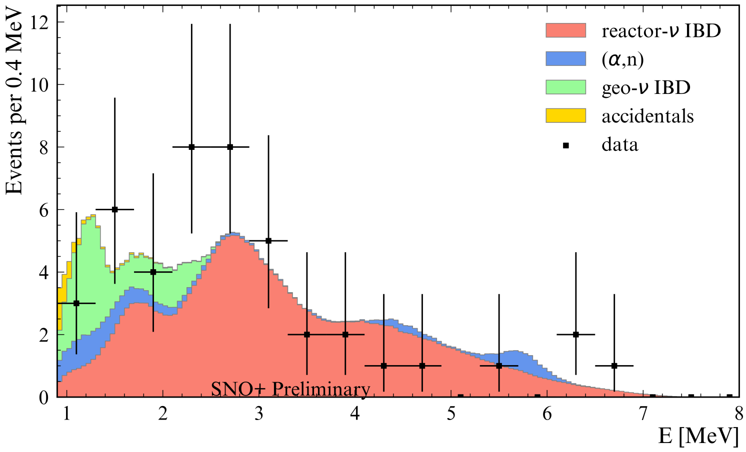

(a) Unconstrained.

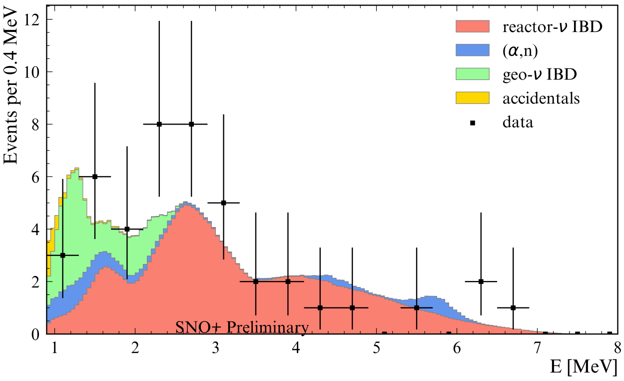

(b) Constrained.

Figure 8.10: Fit prompt energy spectra from oscillation analyses with unconstrained and constrained Δm^2_{21} and s^2_{12}, and where the (α, n)-IBD classifier was applied to both data and PDFs.

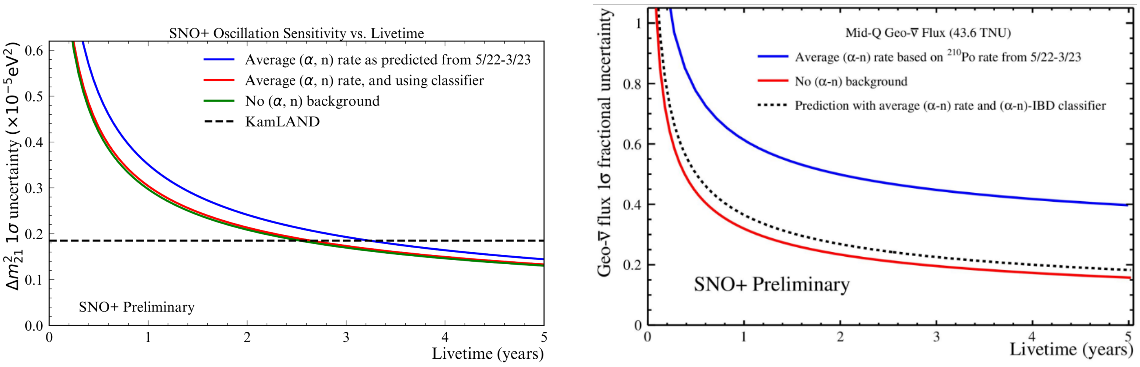

(a) Δm^2_{21} sensitivity. Current best measurement by KamLAND shown for comparison

(b) Geo-ν flux sensitivity, from Anthony Zummo, based on the results of this thesis.

Figure 8.11: Expected measurement sensitivities at SNO+ according to livetime t, based on Azimov datasets at various times subsequently fit to the curve $A/\sqrt{t} + B$. Includes the impact of using the (α, n)-IBD classifier.

Sensitivity Analysis

While the (α, n)-IBD classifier does not yet improve the oscillation analysis due to limited statistics, SNO+ will continue to collect, and has already collected, more data. Therefore, as statistics become less of a limiting factor over time, this classifier is expected to find increasing relevance in improving the sensitivity of SNO+ measurements. To assess how much improvement will be seen over time, similar oscillation analyses were performed with Azimov datasets produced with PDG value oscillation constants and various live-times, and the uncertainty in different fit parameters recorded in each case. These are shown in figure 8.11 for the two parameters of most interest: Δm^2_{21} and the geo-neutrino flux.

According to these simulation-based calculations, using the classifier effectively eliminates (α, n) as a meaningful background to these measurements. It reduces the amount of time needed to achieve a certain precision, and increases the overall achievable precision. The impact is particularly flagrant for the geo-ν flux, as one might expect from the geo-neutrinos and proton recoils being essentially on-top of each-other in the prompt energy spectrum. Meanwhile, the Δm^2_{21} measurement appears in reach of being the world's most precise in the near future, surpassing that of KamLAND. It is hoped that these predictions will be borne out by the data as its quantity increases.

8.5 Summary

The signal and background models described in detail in previous chapters were couched in the context of a numerical log-likelihood maximisation algorithm, designed to fit the long baseline oscillation parameters from reactor antineutrinos. A fast code package was developed, incorporating all the systematic uncertainties previously discussed, including energy scaling and smearing, and including the matter effects algorithm in its entirety. The 57 data points tagged from the 134.5 days of SNO+ livetime in chapter 6 were then passed through this framework.

The oscillation parameter measurements were found to be in agreement with PDG values in all cases, and that of Δm^2_{21} is at present the second most precise in the world, behind KamLAND's [114]. Meanwhile, a geo-neutrino flux is tentatively measured at SNO+, the third location worldwide, as presented at the Neutrino2024 conference in Milan. The (α, n)-IBD classifier developed in this work was found to not improve the accuracy of these results due to the limited sample size, but otherwise agreed with them. To summarise, the final best results from this analysis alone are

$$\Delta m^2_{21} = \left(8.06^{+0.50}_{-0.44}\right) \times 10^{-5}\ \mathrm{eV}^2, \quad s^2_{12} = 0.43^{+0.19}_{-0.33}, \quad \Phi = 70^{+44}_{-40}\ \mathrm{TNU}. \tag{8.42}$$

In future analyses, the classifier is poised to play an important role in improving these measurements, as the analysis becomes less statistics dominated. In part thanks to this, Δm^2_{21} holds the promise of becoming competitive in the next year or two. The geo-neutrino flux measurement promises even greater improvement over time thanks to this classifier, from today's preliminary measurement to the hopeful evidence of tomorrow.

Further improvements to the analysis methodology can be implemented as the data increases. For example, event times can be accounted for: the signal and background fluxes are expected to vary differently over time. Therefore a 2-D PDF in time and prompt energy $f(E,t) = f_E(E)f_t(t)$ would be fit, instead of just a 1D PDF in energy $f_E(E)$ as in the present analysis. Otherwise, the reactor IBD flux's systematic uncertainty could be split into what affects the whole flux together and what affects each reactor individually. These latter reactor uncertainties could easily be floated independently and fit for a few of the closest reactor complexes. Lastly, uncertainties can be further studied and constrained, particularly relating to the (α, n)-IBD classifier. These and other improvements will produce ever greater benefits over time.

Chapter 9

Conclusions

> "Well, here at last, dear friends, on the shores of the Sea comes the end of our fellowship in Middle-earth. Go in peace!"
>
> Gandalf – *The Lord of the Rings*, J.R.R. Tolkien

A detailed calculation of the expected reactor antineutrino induced IBD flux at SNO+ was performed, using new reactor-based antineutrino emission spectra. This was followed by the use of a novel accurate and fast 3-flavour neutrino oscillation algorithm that accounts for the effects of crossing matter with constant density. An overview of the relevant backgrounds and strategies to remove them followed, particularly the troublesome (α, n) events. A model with appropriate systematic uncertainties was thus developed, and implemented in a purpose-made oscillation analysis framework.

57 tagged events from 134.5 days of SNO+ data were analysed in this fashion, yielding measurements of Δm^2_{21} and θ_{12} compatible with the current best constraints, and boasting the second best precision in the case of Δm^2_{21}. A preliminary measurement of the geo-neutrino flux at SNO+ was also produced. Using a new (α, n)-IBD classifier to further purify the data led to consistent oscillation analysis results, though with no improvements at this statistics limited stage.

Potential near-term improvements to this analysis were discussed, to be developed in tandem with increased data collection. The remaining uncertainties associated with the classifier are expected to be reduced in the future, with internal deployments of the AmBe source allowing better calibration of proton recoil simulations. The impact of the classifier on the oscillation analysis will also likely become more important as more data is collected. It is expected to contribute to a potentially world-leading measurement of Δm^2_{21} in the next couple of years, and increasingly accurate measurements of the geo-neutrino flux. Further study of (α, n) interactions themselves may also be on the table.

Overall, this thesis is presented in the context of ongoing analyses on the presented topics, with continuous incremental improvements. Some of these were not able to make it into the present work. As such, the future of these analyses appears very promising. The author aspires for the analysis framework and classifier developed here to be some part of these upcoming accomplishments.

Bibliography

[1] Charlie Mills. *Improved sensitivity to Δm^2_{21} by classification of the $^{13}C(\alpha, n)^{16}O$ background in the SNO+ antineutrino analysis.* PhD thesis, University of Sussex, 2022.

[2] 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.

[3] Anthony Zummo. *Measurement of reactor antineutrino oscillation with SNO+.* PhD thesis, University of Pennsylvania, 2024.

[4] Carlo Giunti and Chung W Kim. *Fundamentals of neutrino physics and astrophysics.* Oxford university press, 2007.

[5] Frederick Reines, Clyde L Cowan Jr, FB Harrison, AD McGuire, and HW Kruse. Detection of the free antineutrino. *Physical Review*, 117(1):159, 1960.

[6] R. L. Workman and Others. *Review of Particle Physics*, volume 2022. 2022.

[7] V Albanese, R Alves, MR Anderson, S Andringa, L Anselmo, E Arushanova, S Asahi, M Askins, DJ Auty, AR Back, et al. The sno+ experiment. *Journal of Instrumentation*, 16(08):P08059, 2021.

[8] SNO+ Collaboration, A Allega, MR Anderson, S Andringa, J Antunes, M Askins, DJ Auty, A Bacon, N Barros, F Barao, et al. Observation of antineutrinos from distant reactors using pure water at sno+. *arXiv preprint arXiv:2210.14154*, 2022.

[9] SNO+ Collaboration, A Allega, MR Anderson, S Andringa, M Askins, DJ Auty, A Bacon, J Baker, F Barão, N Barros, et al. Initial measurement of reactor antineutrino oscillation at sno+. *arXiv preprint arXiv:2405.19700*, 2024.

[10] James Page. Fast exact algorithm for neutrino oscillation in constant matter density. *Computer Physics Communications*, 300:109200, 2024.

[11] Mark Thomson. *Modern particle physics*, pages 329–336. Cambridge University Press, sixth edition, 2013.

[12] Y. Fukuda, T. Hayakawa, E. Ichihara, K. Inoue, K. Ishihara, H. Ishino, Y. Itow, T. Kajita, J. Kameda, S. Kasuga, and et al. Evidence for oscillation of atmospheric neutrinos. *Physical Review Letters*, 81(8):1562–1567, Aug 1998.

[13] Ettore Majorana and Luciano Maiani. A symmetric theory of electrons and positrons. In *Ettore Majorana Scientific Papers*, pages 201–233. Springer, 2006.

[14] C Itzykson and JB Zuber. International series in pure and applied physics. *Quantum field theory*, 1980.

[15] Vedran Brdar, Alexander J Helmboldt, Sho Iwamoto, and Kai Schmitz. Type 1 seesaw mechanism as the common origin of neutrino mass, baryon asymmetry, and the electroweak scale. *Physical Review D*, 100(7):075029, 2019.

[16] Mikael Beuthe. Oscillations of neutrinos and mesons in quantum field theory. *Physics Reports*, 375(2-3):105–218, 2003.

[17] Marek Zralek. From kaons to neutrinos: Quantum mechanics of particle oscillations. *arXiv preprint hep-ph/9810543*, 1998.

[18] Boris Kayser. On the quantum mechanics of neutrino oscillation. *Physical Review D*, 24(1):110, 1981.

[19] Shmuel Nussinov. Solar neutrinos and neutrino mixing. *Physics Letters B*, 63(2):201–203, 1976.

[20] Jack Dunger. *Topological and time based event classification for neutrinoless double beta decay in liquid scintillator*. PhD thesis, University of Oxford, 2018.

[21] KA Olive. Review of particle physics. *Chinese Physics. C, High Energy Physics and Nuclear Physics*, 40(FERMILAB-PUB-16-568), 2016.

[22] Tanner Kaptanoglu. *Reactor Antineutrinos in the SNO+ Water Phase and Detector R&D for Large-Scale Neutrino Detectors*. PhD thesis, University of Pennsylvania, 2020.

[23] J Boger, RL Hahn, JK Rowley, AL Carter, B Hollebone, D Kessler, I Blevis, F Dalnoki-Veress, A DeKok, J Farine, et al. 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.

[24] Q. R. Ahmad, R. C. Allen, T. C. Andersen, J. D.Anglin, J. C. Barton, E. W. Beier, M. Bercovitch, J. Bigu, and others. Direct evidence for neutrino flavor transformation from neutral-current interactions in the sudbury neutrino observatory. *Phys. Rev. Lett.*, 89:011301, Jun 2002.

[25] G Alimonti, C Arpesella, MB Avanzini, H Back, M Balata, D Bartolomei, A De Bellefon, G Bellini, J Benziger, A Bevilacqua, et al. The liquid handling systems for the borexino solar neutrino detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 609(1):58–78, 2009.

[26] J Benziger, L Cadonati, Frank Calaprice, M Chen, A Corsi, F Dalnoki-Veress, R Fernholz, R Ford, Cristiano Galbiati, A Goretti, et al. A scintillator purification system for the borexino solar neutrino detector. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 587(2-3):277–291, 2008.

[27] K Eguchi et al. and kamland collaboration 2003 first results from kamland: evidence for reactor anti-neutrino disappearance. *Phys. Rev. Lett*, 90:021802, 2003.

[28] G Alimonti, C Arpesella, H Back, M Balata, D Bartolomei, A De Bellefon, G Bellini, J Benziger, A Bevilacqua, D Bondi, et al. The borexino detector at the laboratori nazionali del gran sasso. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 600(3):568–593, 2009.

[29] CUORE collaboration et al. Measurement of the $2\nu\beta\beta$ decay half-life of 130te with cuore. *Phys. Rev. Lett*, 126:171801, 2021.

[30] SNO+ Collaboration, MR Anderson, S Andringa, M Askins, DJ Auty, F Barão, N Barros, R Bayes, EW Beier, A Bialek, et al. Optical calibration of the sno+ detector in the water phase with deployed sources. *arXiv preprint arXiv:2106.03951*, 2021.

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

[32] S Hans, R Rosero, L Hu, O Chkvorets, WT Chan, S Guan, W Beriguete, A Wright, R Ford, MC Chen, et al. Purification of telluric acid for sno+ neutrinoless double-beta decay search. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 795:132–139, 2015.

[33] T. Bolton. The braidwood reactor anitneutrino experiment. *Nuclear Physics B - Proceedings Supplements*, 149:166–169, 2005. NuFact04.

[34] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, and others. Geant4—a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003.

[35] Elisabeth Falk, Jeffrey Lidgard, Mark I Stringer, and Esther Turner. Commissioning of ellie for sno+. *arXiv preprint arXiv:1705.00354*, 2017.

[36] SNO+ Collaboration, A Allega, MR Anderson, S Andringa, M Askins, DJ Auty, A Bacon, J Baker, F Barão, N Barros, et al. Measurement of the ^{8} b solar neutrino flux using the full sno+ water phase. *arXiv preprint arXiv:2407.17595*, 2024.

[37] Anil Kumar and Sanjib Kumar Agarwalla. Validating the earth's core using atmospheric neutrinos with ical at ino. *arXiv preprint arXiv:2104.11740*, 2021.

[38] Frederick Reines. The neutrino: From poltergeist to particle. *Reviews of Modern Physics*, 68(2):317, 1996.

[39] FP An, AB Balantekin, HR Band, M Bishai, S Blyth, D Cao, GF Cao, J Cao, WR Cen, YL Chan, et al. New measurement of θ 13 via neutron capture on hydrogen at daya bay. *Physical Review D*, 93(7):072011, 2016.

[40] Feng Peng An, Akif Baha Balantekin, Henry Reyer Band, M Bishai, Simon Blyth, D Cao, Guo Fu Cao, Jun Cao, WR Cen, Yat Long Chan, et al. Improved measurement of the reactor antineutrino flux and spectrum at daya bay. *Chinese Physics C*, 41(1):013002, 2017.

[41] X. B. Ma, W. L. Zhong, L. Z. Wang, Y. X. Chen, and J. Cao. Improved calculation of the energy release in neutron-induced fission. *Phys. Rev. C*, 88:014605, Jul 2013.

[42] Th A Mueller, D Lhuillier, Muriel Fallot, A Letourneau, S Cormon, M Fechner, Lydie Giot, Th Lasserre, J Martino, G Mention, et al. Improved predictions of reactor antineutrino spectra. *Physical Review C*, 83(5):054615, 2011.

[43] IAEA. Operating experience with nuclear power stations in member states, 2023.

[44] World Nuclear Association. Nuclear power reactors. `https://world-nuclear.org/information-library/nuclear-fuel-cycle/nuclear-power-reactors/nuclear-power-reactors`. Accessed: 2024-07-15.

[45] Daya Bay Collaboration et al. Extraction of the 235u and 239pu antineutrino spectra at daya bay. *Journal of Environmental Sciences (China) English Ed*, 2019.

[46] Patrick Huber. Determination of antineutrino spectra from nuclear reactors. *Physical Review C*, 84(2):024617, 2011.

[47] Patrick Huber. Reactor antineutrino fluxes–status and challenges. *Nuclear Physics B*, 908:268–278, 2016.

[48] Pouya Bakhti, Min-Gwa Park, Meshkat Rajaee, Chang Sub Shin, and Seodong Shin. Revisiting reactor anti-neutrino 5 mev bump with ^{13}c neutral-current interaction. *arXiv preprint arXiv:2405.08724*, 2024.

[49] David Kaleko. *MicroBooNE: The Search For The MiniBooNE Low Energy Excess*. Columbia University, 2017.

[50] FP An, AB Balantekin, M Bishai, S Blyth, GF Cao, J Cao, JF Chang, Y Chang, HS Chen, SM Chen, et al. Antineutrino energy spectrum unfolding based on the daya bay measurement and its applications. *Chinese Physics C*, 45(7):073001, 2021.

[51] FP An, M Andriamirado, AB Balantekin, HR Band, CD Bass, DE Bergeron, D Berish, M Bishai, S Blyth, NS Bowden, et al. Joint determination of reactor antineutrino spectra from u 235 and pu 239 fission by daya bay and prospect. *Physical review letters*, 128(8):081801, 2022.

[52] *Operating Experience with Nuclear Power Stations in Member States*. Operating Experience with Nuclear Power Stations in Member States. INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2022.

[53] *Operating Experience with Nuclear Power Stations in Member States 2023 Edition*. Operating Experience with Nuclear Power Stations in Member States. INTERNATIONAL ATOMIC ENERGY AGENCY, Vienna, 2023.

[54] IESO. Ieso monthly electrical data. http://reports.ieso.ca/public/GenOutputCapabilityMonth/. Accessed: 2024-08-5.

[55] Giulia Ricciardi, Natascia Vignaroli, and Francesco Vissani. A discussion of the cross section $\bar{nu}_e + p \to e^+ + n$. *arXiv e-prints*, pages arXiv–2311, 2023.

[56] Alessandro Strumia and Francesco Vissani. Precise quasielastic neutrino/nucleon cross-section. *Physics Letters B*, 564(1-2):42–54, 2003.

[57] P Vogel and JF Beacom. The angular distribution of the reaction νe+ p$\to$ e+ n. *Phys. Rev. D*, 60:053003, 1999.

[58] D.H. Wilkinson. Phase space for neutron beta-decay: an update. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 404(2):305–310, 1998.

[59] J. C. Hardy and I. S. Towner. Superallowed fermi beta decay and coulomb mixing in nuclei. In *AIP Conference Proceedings*. AIP, 1999.

[60] Chien-Yeah Seng, Mikhail Gorchtein, and Michael J. Ramsey-Musolf. Dispersive evaluation of the inner radiative correction in neutron and nuclear <mml:math xmlns:mml="http://www.w3.org/1998/math/mathml" display="inline"><mml:mi>β</mml:mi></mml:math> decay. *Physical Review D*, 100(1), July 2019.

[61] Particle Data Group, P A Zyla, R M Barnett, J Beringer, O Dahl, D A Dwyer, D E Groom, C J Lin, K S Lugovsky, E Pianori, D J Robinson, C G Wohl, W M Yao, K Agashe, G Aielli, B C Allanach, C Amsler, M Antonelli, E C Aschenauer, D M Asner, H Baer, Sw Banerjee, L Baudis, C W Bauer, J J Beatty, V I Belousov, S Bethke, A Bettini, O Biebel, K M Black, E Blucher, O Buchmuller, V Burkert, M A Bychkov, R N Cahn, M Carena, A Ceccucci, A Cerri, D Chakraborty, R Sekhar Chivukula, G Cowan, G D'Ambrosio, T Damour, D de Florian, A de Gouvêa, T DeGrand, P de Jong, G Dissertori, B A Dobrescu, M D'Onofrio, M Doser, M Drees, H K Dreiner, P Eerola, U Egede, S Eidelman, J Ellis, J Erler, V V Ezhela, W Fetscher, B D Fields, B Foster, A Freitas, H Gallagher, L Garren, H J Gerber, G Gerbier, T Gershon, Y Gershtein, T Gherghetta, A A Godizov, M C Gonzalez-Garcia, M Goodman, C Grab, A V Gritsan, C Grojean, M Grünewald, A Gurtu, T Gutsche, H E Haber, C Hanhart, S Hashimoto, Y Hayato, A Hebecker, S Heinemeyer, B Heltsley, J J Hernández-Rey, K Hikasa, J Hisano, A Höcker, J Holder, A Holtkamp, J Huston, T Hyodo, K F Johnson, M Kado, M Karliner, U F Katz, M Kenzie, V A Khoze, S R Klein, E Klempt, R V Kowalewski, F Krauss, M Kreps, B Krusche, Y Kwon, O Lahav, J Laiho, L P Lellouch, J Lesgourgues, A R Liddle, Z Ligeti, C Lippmann, T M Liss, L Littenberg, C Lourengo, S B Lugovsky, A Lusiani, Y Makida, F Maltoni, T Mannel, A V Manohar, W J Marciano, A Masoni, J Matthews, U G Meißner, M Mikhasenko, D J Miller, D Milstead, R E Mitchell, K Mönig, P Molaro, F Moortgat, M Moskovic,

K Nakamura, M Narain, P Nason, S Navas, M Neubert, P Nevski, Y Nir, K A Olive, C Patrignani, J A Peacock, S T Petcov, V A Petrov, A Pich, A Piepke, A Pomarol, S Profumo, A Quadt, K Rabbertz, J Rademacker, G Raffelt, H Ramani, M Ramsey-Musolf, B N Ratcliff, P Richardson, A Ringwald, S Roesler, S Rolli, A Romaniouk, L J Rosenberg, J L Rosner, G Rybka, M Ryskin, R A Ryutin, Y Sakai, G P Salam, S Sarkar, F Sauli, O Schneider, K Scholberg, A J Schwartz, J Schwiening, D Scott, V Sharma, S R Sharpe, T Shutt, M Silari, T Sjöstrand, P Skands, T Skwarnicki, G F Smoot, A Soffer, M S Sozzi, S Spanier, C Spiering, A Stahl, S L Stone, Y Sumino, T Sumiyoshi, M J Syphers, F Takahashi, M Tanabashi, J Tanaka, M Taševský, K Terashi, J Terning, U Thoma, R S Thorne, L Tiator, M Titov, N P Tkachenko, D R Tovey, K Trabelsi, P Urquijo, G Valencia, R Van de Water, N Varelas, G Venanzoni, L Verde, M G Vincter, P Vogel, W Vogelsang, A Vogt, V Vorobyev, S P Wakely, W Walkowiak, C W Walter, D Wands, M O Wascko, D H Weinberg, E J Weinberg, M White, L R Wiencke, S Willocq, C L Woody, R L Workman, M Yokoyama, R Yoshida, G Zanderighi, G P Zeller, O V Zenin, R Y Zhu, S L Zhu, F Zimmermann, J Anderson, T Basaglia, V S Lugovsky, P Schaffner, and W Zheng. Review of Particle Physics. *Progress of Theoretical and Experimental Physics*, 2020(8):083C01, 08 2020.

[62] Bastian Märkisch, Holger Mest, Heiko Saul, Xiangzun Wang, Hartmut Abele, Dirk Dubbers, Michael Klopf, Alexander Petoukhov, Christoph Roick, Torsten Soldner, et al. Measurement of the weak axial-vector coupling constant in the decay of free neutrons using a pulsed cold neutron beam. *Physical review letters*, 122(24):242501, 2019.

[63] Benjamin Tam. Sno+ scintillator (lab + 2.2 g/l ppo) density plot, sno+ internal document 7405. Accessed: 2024-07-010.

[64] Benjamin Sum Ki Tam. *Enabling Neutrinoless Double Beta Decay in the SNO+ Experiment through the Deployment and Study of Liquid Scintillator*. PhD thesis, Queen's U., Kingston, 2023.

[65] CEPSA. Lab certificate of analysis, sno+ internal document 7572. Accessed: 2024-07-011.

[66] National institute of standards and technology. physics measurement laboratory. hydrogen. https://physics.nist.gov/cgi-bin/Compositions/stand_alone.pl?ele=H. Accessed: 2024-07-011.

[67] Position recoordination for bis-msb, sno+ internal document 8326. https://www.snolab.ca/snoplus/private/DocDB/0083/008326/002/BisMSBPosRecon_20240717.pdf. Accessed: 2024-09-19.

[68] HW Zaglauer and KH Schwarzer. The mixing angles in matter for three generations of neutrinos and the msw mechanism. *Zeitschrift für Physik C Particles and Fields*, 40(2):273–282, 1988.

[69] James P Kneller and Gail C McLaughlin. Three flavor neutrino oscillations in matter: Flavor diagonal potentials, the adiabatic basis, and the c p phase. *Physical Review D*, 80(5):053002, 2009.

[70] Patrick Huber, Joachim Kopp, Manfred Lindner, Mark Rolinec, and Walter Winter. Globes: General long baseline experiment simulator. *Computer physics communications*, 177(5):439–440, 2007.

[71] Joachim Kopp. Efficient numerical diagonalization of hermitian 3×3 matrices. *International Journal of Modern Physics C*, 19(03):523–548, 2008.

[72] V Barger, K Whisnant, S Pakvasa, and RJN Phillips. Matter effects on three-neutrino oscillations. *Physical Review D*, 22(11):2718, 1980.

[73] Tommy Ohlsson and Håkan Snellman. Neutrino oscillations with three flavors in matter: Applications to neutrinos traversing the earth. *Physics Letters B*, 474(1-2):153–162, 2000.

[74] Peter B Denton, Stephen J Parke, and Xining Zhang. Neutrino oscillations in matter via eigenvalues. *Phys. Rev. D*, 101(9):093001, 2020.

[75] Gabriela Barenboim, Peter B Denton, Stephen J Parke, and Christoph A Ternes. Neutrino oscillation probabilities through the looking glass. *Physics Letters B*, 791:351–360, 2019.

[76] Chee Sheng Fong. Analytic neutrino oscillation probabilities. *SciPost Physics*, 15(1):013, 2023.

[77] Bushra Shafaq and Faisal Akram. The effect of the earth matter on three neutrino oscillations and sensitivity to cp phase parameter. *The European Physical Journal Plus*, 135(1):94, 2020.

[78] Mauricio Bustamante. Nuoscprobexact: a general-purpose code to compute exact two-flavor and three-flavor neutrino oscillation probabilities. *arXiv preprint arXiv:1904.12391*, 2019.

[79] Walter D Mooney and Mikhail K Kaban. The north american upper mantle: Density, composition, and evolution. *Journal of Geophysical Research: Solid Earth*, 115(B12), 2010.

[80] Don L Anderson. *Theory of the Earth*. Blackwell scientific publications, 1989.

[81] NOAA. National centers for environmental information (https://www.ngdc.noaa.gov/), Feb 2012.

[82] A Piepke, Palo Verde Collaboration, et al. Final results from the palo verde neutrino oscillation experiment. *Progress in Particle and Nuclear Physics*, 48(1):113–121, 2002.

[83] Marco Apollonio, A Baldini, C Bemporad, E Caffau, Fabrizio Cei, Y Declais, H De Kerret, B Dieterle, A Etenko, J George, et al. Limits on neutrino oscillations from the chooz experiment. *Physics Letters B*, 466(2-4):415–430, 1999.

[84] H. de Kerret et al. Double Chooz θ_{13} measurement via total neutron capture detection. *Nature Phys.*, 16(5):558–564, 2020.

[85] G Bak, JH Choi, HI Jang, JS Jang, SH Jeon, KK Joo, Kiwon Ju, DE Jung, JG Kim, JH Kim, et al. Measurement of reactor antineutrino oscillation amplitude and frequency at reno. *Physical Review Letters*, 121(20):201801, 2018.

[86] FP An, WD Bai, AB Balantekin, M Bishai, S Blyth, GF Cao, J Cao, JF Chang, Y Chang, HS Chen, et al. Precision measurement of reactor antineutrino oscillation at kilometer-scale baselines by daya bay. *Physical review letters*, 130(16):161802, 2023.

[87] B Aharmim, SN Ahmed, AE Anthony, N Barros, EW Beier, Alain Bellerive, B Beltran, M Bergevin, SD Biller, K Boudjemline, et al. Combined analysis of all three phases of solar neutrino data from the sudbury neutrino observatory. *Physical Review C—Nuclear Physics*, 88(2):025501, 2013.

[88] Bruce T Cleveland, Timothy Daily, Raymond Davis Jr, James R Distel, Kenneth Lande, CK 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, 1998.

[89] Kohji S Hirata, Takaaki Kajita, T Kifune, K Kihara, Masayuki Nakahata, K Nakamura, S Ohara, Y Oyama, N Sato, M Takita, et al. Observation of b 8 solar neutrinos in the kamiokande-ii detector. *Physical review letters*, 63(1):16, 1989.

[90] Michael Altmann, M Balata, P Belli, E Bellotti, R Bernabei, E Burkert, C Cattadori, R Cerulli, Marco Chiarini, M Cribier, et al. Complete results for five years of gno solar neutrino observations. *Physics Letters B*, 616(3-4):174–190, 2005.

[91] Vladimir N Gavrin, JN Abdurashitovi, TJ Bowles, ML Cherry, BT Cleveland, R Davis Jr, SR Elliott, SV Girin, VV Gorbachev, TV Ibragimova, et al. Solar neutrino results from sage. *Nuclear Physics B-Proceedings Supplements*, 91(1-3):36–43, 2001.

[92] Gabriel D Orebi Gann, Kai Zuber, Daniel Bemmerer, and Aldo Serenelli. The future of solar neutrinos. *Annual Review of Nuclear and Particle Science*, 71(1):491–528, 2021.

[93] Zihao Bo, Wei Chen, Xun Chen, Yunhua Chen, Zhaokan Cheng, Xiangyi Cui, Yingjie Fan, Deqing Fang, Zhixing Gao, Lisheng Geng, et al. First indication of solar ^{8} b neutrino flux through coherent elastic neutrino-nucleus scattering in pandax-4t. *arXiv preprint arXiv:2407.10892*, 2024.

[94] E Aprile, J Aalbers, K Abe, S Ahmed Maouloud, L Althueser, B Andrieu, E Angelino, D Antón Martin, F Arneodo, L Baudis, et al. First measurement of solar ^{8} b neutrinos via coherent elastic neutrino-nucleus scattering with xenonnt. *arXiv preprint arXiv:2408.02877*, 2024.

[95] K Abe, C Bronner, Y Hayato, K Hiraide, K Hosokawa, K Ieki, M Ikeda, S Imaizumi, K Iyogi, J Kameda, et al. Solar neutrino measurements using the full data period of super-kamiokande-iv. *Physical Review D*, 109(9):092001, 2024.

[96] Marco Pallavicini et al. (Borexino collaboration). Solar neutrino results from borexino and main future perspectives. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 630(1):210–213, 2011. Proceedings of the 2nd Roma International Conference on Astroparticle Physics (RICAP 2009).

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

[98] KamLAND Collaboration. K. eguchi et al., first results from kamland: Evidence for reactor anti-neutrino disappearance. *Phys. Rev. Lett*, 90:021802, 2003.

[99] Azusa Gando, Y Gando, H Hanakago, Hirokazu Ikeda, K Inoue, K Ishidoshiro, H Ishikawa, M Koga, R Matsuda, S Matsuda, et al. Reactor on-off antineutrino measurement with kamland. *Physical Review D—Particles, Fields, Gravitation, and Cosmology*, 88(3):033001, 2013.

[100] Peter B Denton and Stephen J Parke. Fast and accurate algorithm for calculating long-baseline neutrino oscillation probabilities with matter effects: Nufast. *arXiv preprint arXiv:2405.02400*, 2024.

[101] Gcc: Options that control optimization. `https://gcc.gnu.org/onlinedocs/gcc/Optimize-Options.html#:~:text=%2DOfast%20enables%20all%20%2DO3%20optimizations,%2Dfno%2Dprotect%2Dparens%20.` Accessed: 2024-12-09.

[102] Simon Byrne. Simon's notes: Beware of fast-math. `https://simonbyrne.github.io/notes/fastmath/`. Accessed: 2024-12-09.

[103] CL Cowan Jr and F Reines. 19. detection of the free neutrino: a confirmation. cl cowan. *Neutrinos and Other Matters: Selected Works of Frederick Reines*, 1990.

[104] File:decay chain(4n+2, uranium series).svg. `https://commons.wikimedia.org/wiki/File:Decay_chain(4n%2B2,_Uranium_series).svg`. Accessed: 2024-07-02.

[105] File:decay chain of thorium.svg. `https://commons.wikimedia.org/wiki/File:Decay_Chain_of_Thorium.svg`. Accessed: 2024-07-02.

[106] SNO+ Collaboration, A Allega, MR Anderson, S Andringa, M Askins, DJ Auty, A Bacon, J Baker, F Barão, N Barros, et al. Initial measurement of reactor antineutrino oscillation at sno+. *arXiv preprint arXiv:2405.19700*, 2024.

[107] Takaaki Kajita. Atmospheric neutrinos. *Advances in High Energy Physics*, 2012(1):504715, 2012.

[108] Antoine Kouchner. Next-generation atmospheric neutrino experiments. *Physics of the Dark Universe*, 4:60–74, 2014.

[109] Costas Andreopoulos, A Bell, D Bhattacharya, F Cavanna, J Dobson, S Dytman, H Gallagher, P Guzowski, R Hatcher, P Kehayias, et al. The genie neutrino monte carlo generator. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 614(1):87–104, 2010.

[110] Cindy Lin. Atmospheric neutrinos in partial fill, sno+ internal document 6958. Accessed: 2024-07-011.

[111] Bassam Aharmim, SN Ahmed, TC Andersen, AE Anthony, N Barros, EW Beier, Alain Bellerive, B Beltran, M Bergevin, SD Biller, et al. Measurement of the cosmic ray and neutrino-induced muon flux at the sudbury neutrino observatory. *Physical Review D*, 80(1):012001, 2009.

[112] Billy Liggins. *Cosmic muon-induced neutrons in the SNO+ water phase*. PhD thesis, Queen Mary University of London, 2020.

[113] S Abe, S Enomoto, K Furuno, Y Gando, H Ikeda, K Inoue, Y Kibe, Y Kishimoto, M Koga, Y Minekawa, et al. Production of radioactive isotopes through cosmic muon spallation in kamland. *Physical Review C—Nuclear Physics*, 81(2):025807, 2010.

[114] Shingo Abe, T Ebihara, S Enomoto, K Furuno, Y Gando, K Ichimura, H Ikeda, K Inoue, Y Kibe, Y Kishimoto, et al. Precision measurement of neutrino oscillation parameters with kamland. *Physical review letters*, 100(22):221803, 2008.

[115] M Agostini, K Altenmüller, S Appel, V Atroshchenko, Z Bagdasarian, D Basilico, G Bellini, J Benziger, D Bick, G Bonfini, et al. Comprehensive geoneutrino analysis with borexino. *Physical Review D*, 101(1):012009, 2020.

[116] Shirley Weishi Li and John F Beacom. First calculation of cosmic-ray muon spallation backgrounds for mev astrophysical neutrino signals in super-kamiokande. *Physical Review C*, 89(4):045801, 2014.

[117] Ricardo Arevalo Jr, William F McDonough, and Mario Luong. The k/u ratio of the silicate earth: Insights into mantle composition, structure and thermal evolution. *Earth and Planetary Science Letters*, 278(3-4):361–369, 2009.

[118] Ingrida Semenec. *The Geoneutrino Signal in the SNO+ experiment*. PhD thesis, 2023.

[119] Geoneutrino spectrum and luminosity. `https://www.awa.tohoku.ac.jp/~sanshiro/research/geoneutrino/spectrum/`. Accessed: 2024-08-18.

[120] Enomoto Sanshiro. *Neutrino Geophysics and Observation of Geo-Neutrinos at KamLAND,*. PhD thesis, Tohoku University, 2005.

[121] M Agostini, K Altenmüller, S Appel, V Atroshchenko, Z Bagdasarian, D Basilico, G Bellini, J Benziger, D Bick, G Bonfini, et al. Comprehensive geoneutrino analysis with borexino. *Physical Review D*, 101(1):012009, 2020.

[122] S Abe, S Asami, M Eizuka, S Futagi, A Gando, Y Gando, T Gima, A Goto, T Hachiya, K Hata, et al. Abundances of uranium and thorium elements in earth estimated by geoneutrino spectroscopy. *Geophysical Research Letters*, 49(16):e2022GL099566, 2022.

[123] Ondřej Šrámek, Bedřich Roskovec, Scott A Wipperfurth, Yufei Xi, and William F McDonough. Revealing the earth's mantle from the tallest mountains using the jinping neutrino experiment. *Scientific Reports*, 6(1):33034, 2016.

[124] Neutron sources and backgrounds in sno+, sno+ internal document 2497. Accessed: 2024-07-02.

[125] Gilles Delaygue. *Oxygen Isotopes*, pages 666–673. Springer Netherlands, Dordrecht, 2009.

[126] Www table of radioactive isotopes, ^{210}po. http://nucleardata.nuclear.lu.se/toi/nuclide.asp?iZA=840210. Accessed: 2024-07-02.

[127] March 2023 collaboration meeting: Major background highlights, sno+ internal document 7712. Accessed: 2024-07-02.

[128] Toru Murata, Hiroyuki Matsunobu, and Keiichi Shibata. Evaluation of the (α, xn) reaction data for jendl/an-2005. Technical report, Japan Atomic Energy Agency, 2006.

[129] KK Sekharan, AS Divatia, MK Mehta, SS Kerekatte, and KB Nambiar. *sup*13 c (*alpha*, n) *sup*16 o reaction cross section between 1.95 and 5.57 mev. Technical report, Atomic Energy Establishment, Trombay, India, 1967.

[130] JK Bair and FX Haas. Total neutron yield from the reactions c 13 (α, n) o 16 and o 1 7, 1 8 (α, n) ne 2 0, 2 1. *Physical Review C*, 7(4):1356, 1973.

[131] S Harissopulos, HW Becker, JW Hammer, A Lagoyannis, C Rolfs, and F Strieder. Cross section of the c 13 (α, n) o 16 reaction: A background for the measurement of geo-neutrinos. *Physical Review C—Nuclear Physics*, 72(6):062801, 2005.

[132] Interactions of ions with matter. http://www.srim.org/. Accessed: 2024-07-02.

[133] Janisweb α-induced cross sections. https://www.oecd-nea.org/janisweb/tree/A. Accessed: 2024-07-07. Used the JENDL/2005 and EXFOR branches.

[134] National institute of standards and technology. physics measurement laboratory. carbon. https://physics.nist.gov/cgi-bin/Compositions/stand_alone.pl?ele=C. Accessed: 2024-07-08.

[135] Note on external backgrounds in full-fill, sno+ internal document 7687. Accessed: 2024-07-09.

[136] Jie Zhao, Ze-Yuan Yu, Jiang-Lai Liu, Xiao-Bo Li, Fei-Hong Zhang, and Dong-Mei Xia. 13c (α, n) 16o background in a liquid scintillator based neutrino experiment. *Chinese Physics C*, 38(11):116201, 2014.

[137] Martin Freer and Hans Otto Uldall Fynbo. The hoyle state in 12c. *Progress in Particle and Nuclear Physics*, 78:1–23, 2014.

[138] Janis nuclear data. https://www.oecd-nea.org/jcms/pl_39910/janis. Accessed: 2024-07-08.

[139] (α, n) nuclear data evaluations and data needs. https://conferences.iaea.org/event/366/attachments/15587/25866/INDC(NDS)-0836.pdf. Summary Report of the Technical Meeting 8-12 November 2021 (virtual event). Accessed: 2024-07-08.

[140] HM O'Keeffe, E O'sullivan, and MC Chen. Scintillation decay time and pulse shape discrimination in oxygenated and deoxygenated solutions of linear alkylbenzene for the sno+ experiment. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 640(1):119–122, 2011.

[141] B Herman, VC Frohlich, JR Lakowicz, DB Murphy, KR Spring, and MW Davidson. Basic concepts in fluorescence. *Microscopy Resource Center Olympus*, 2003.

[142] Measurement of the dependence of the light yields of lab-based and ej-301 scintillators on electron energy, sno+ internal document 964. Accessed: 2024-08-28.

[143] Rat optics 2.2g/l, sno+ internal document 7571. Accessed: 2024-08-28.

[144] EJ Callaghan, BL Goldblum, JA Brown, TA Laplace, Juan J Manfredi, M Yeh, and GD Orebi Gann. Measurement of proton light yield of water-based liquid scintillator. *The European Physical Journal C*, 83(2):134, 2023.

[145] Mario Schwarz, Sabrina M Franke, Lothar Oberauer, Miriam D Plein, Hans Th J Steiger, and Marc Tippmann. Measurements of the lifetime of orthopositronium in the lab-based liquid scintillator of juno. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 922:64–70, 2019.

[146] Y Abe, JC Dos Anjos, JC Barriere, E Baussan, I Bekman, M Bergevin, TJC Bezerra, L Bezrukov, E Blucher, C Buck, et al. Ortho-positronium observation in the double chooz experiment. *Journal of High Energy Physics*, 2014(10):1–17, 2014.

[147] Gianpaolo Bellini, J Benziger, D Bick, G Bonfini, D Bravo, M Buizza Avanzini, B Caccianiga, L Cadonati, Frank Calaprice, P Cavalcante, et al. Final results of borexino phase-i on low-energy solar neutrino spectroscopy. *Physical Review D*, 89(11):112007, 2014.

[148] S Perasso, Giovanni Consolati, D Franco, C Jollet, A Meregaglia, A Tonazzo, and M Yeh. Measurement of ortho-positronium properties in liquid scintillators. *Journal of Instrumentation*, 9(03):C03028, 2014.

[149] Davide Franco, Giovanni Consolati, and D Trezzi. Positronium signature in organic liquid scintillators for neutrino experiments. *Physical Review C*, 83(1):015504, 2011.

[150] Giovanni Consolati, D Franco, S Hans, C Jollet, A Meregaglia, S Perasso, A Tonazzo, and M Yeh. Characterization of positronium properties in doped liquid scintillators. *Physical Review C*, 88(6):065502, 2013.

[151] Sergey V Stepanov and Vsevolod M Byakov. Electric field effect on positronium formation in liquids. *The Journal of chemical physics*, 116(14):6178–6195, 2002.

[152] Glen Cowan. *Statistical data analysis*. Oxford university press, 1998.

[153] Ronald A Fisher. The use of multiple measurements in taxonomic problems. *Annals of eugenics*, 7(2):179–188, 1936.

[154] R Wilkinson. Applied multivariate statistics: Fisher's linear discriminant rule. `https://rich-d-wilkinson.github.io/MATH3030/8.3-FLDA.html`. Accessed: 2024-08-28.

[155] Kenzo Nakamura. Review of particle physics. *Journal of Physics G: Nuclear and Particle Physics*, 37(7 A), 2010.

[156] S. S. Wilks. The large-sample distribution of the likelihood ratio for testing composite hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, 1938.

[157] Steven D Biller and Scott M Oser. Another look at confidence intervals: Proposal for a more relevant and transparent approach. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 774:103–119, 2015.

[158] Gary J Feldman and Robert D Cousins. Unified approach to the classical statistical analysis of small signals. *Physical review D*, 57(7):3873, 1998.

[159] Roofit data analysis framework. `https://root.cern/manual/roofit/`. Accessed: 2024-09-06.

[160] Fred James. Minuit: function minimization and error analysis reference manual. Technical Report 94.1, CERN, 1998. Accessed 2024-09-08.

[161] ATLAS Statistics Forum. Recommendations for presentation of error bars. `https://www.pp.rhul.ac.uk/~cowan/atlas/ErrorBars.pdf`. Accessed: 2024-09-09.

Appendix A

Appendix

A.1 Reactor IBD Prediction Tables

Addition uncertainty and correlation information used in the reaction antineutrino emission spectra computed in chapter 4.

Table A.1: Correlation between fission fraction variations [40].

	^{235}U	^{238}U	^{239}Pu	^{241}Pu
^{235}U	1	-0.22	-0.53	-0.18
^{238}U	-0.22	1	0.18	0.26
^{239}Pu	-0.53	0.18	1	0.49
^{241}Pu	-0.18	0.26	0.49	1

Table A.2: ^{238}U fit parameter uncertainties and correlation matrix [42].

Uncertainty

	a_0	a_1	a_2	a_3	a_4	a_5
	0.124	0.0586	0.0111	0.00192	0.000284	0.0000286

Correlations

	a_0	a_1	a_2	a_3	a_4	a_5
a_0	1.00	-0.86	0.20	0.30	0.08	-0.27
a_1	-0.86	1.00	-0.58	-0.21	0.04	0.23
a_2	0.20	-0.58	1.00	-0.48	-0.17	0.20
a_3	0.30	-0.21	-0.48	1.00	-0.36	-0.20
a_4	0.08	0.04	-0.17	-0.36	1.00	-0.77
a_5	-0.27	0.23	0.20	-0.20	-0.77	1.00

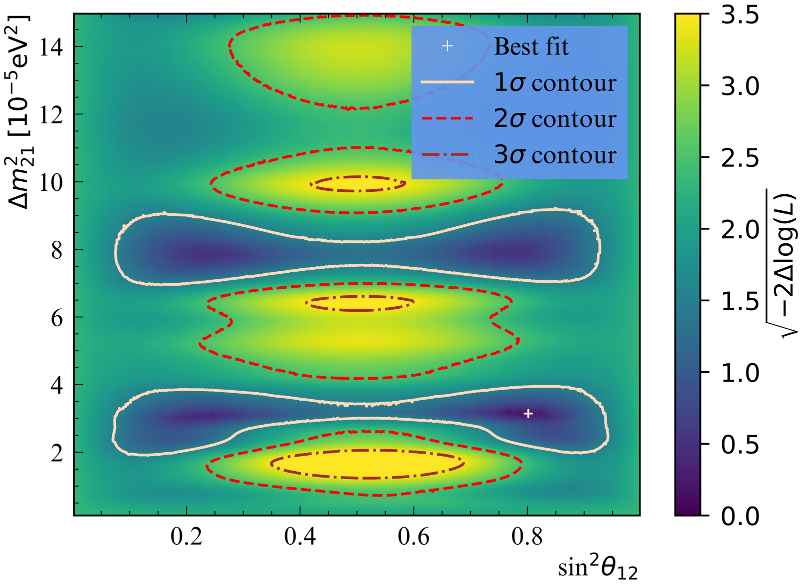

Figure A.1: Oscillation analysis result, when scanning over parameter phase space, and using the (α, n)-IBD classifier. Shows maximised log-likelihood difference phase space, with likelihood contours overlain.

A.2 A Note On Real Data Error Bars

This is a technical note on the displayed error bars for the measured number of data points in a bin. These have no uncertainty, since they show the true number of measured events; it is instead in the models they are compared to that have uncertainty. Now the models have systematic uncertainties based on the fit parameter uncertainties, but also statistical uncertainty simply in the form of Poissonian counting. That is, for each bin, given the estimated expected number of events λ (based on fit parameters), the range of probable measured events n is governed by the Poisson distribution's standard deviation $\sqrt{\lambda}$. However, displaying such model uncertainties on an already busy plot would render it messy at best or illegible at worst. As a way to compromise between legibility and still showing some sense statistical agreement between the data and the models, error bars are added to the sparser data points. These are frequentist uncertainties, encompassing which means λ would give rise to the number of measured points n or more/less (depending on upper of lower bounds) in a combined 68% percent of cases. These are realised using the methods in [161] (equations 6 and 7), themselves taken from [155] (equations 33.59 a and b). In the case of $n = 0$, no error bars are displayed to avoid subtleties on uncertainties in this case. Even for plots that are not busy, and so could easily have model error bands shown, the error bars are added to the data point instead for visual consistency. Such display choices are also fairly common in other publications in the field.

A.3 Additional Oscillation Analysis Plots

A plot that was not displayed in section 8.4.2.

A.4 List of Acronyms

SNO Sudbury Neutrino Observatory

$2\nu\beta\beta$ double beta decay

$0\nu\beta\beta$ neutrinoless double beta decay

SM Standard Model

LH Left Hand

RH Right Hand

NC Neutral Current

CC Charged Current

QFT Quantum Field Theory

CKM Cabbibo-Kobayashi-Maskawa

PMNS Pontecorvo-Maki-Nakagawa-Sakata

V-A Vector-Axial

h.c. Hermitian Conjugate

AV Acrylic Vessel

PMT Photomultiplier Tube

OWL Outward Looking PMT

PSUP PMT Support Structure

PMTIC PMT Interface Card

FEC Front-End Card

DAQ Data Acquisition System

MTC/A+ Analog Master Trigger Card ("+" for SNO+)

DB Daughter Board

GT Global Trigger

GTID Global Trigger ID

PE Photoelectron

TeA Telluric Acid

BD Butane Diol

TeBD Tellurium Butane Diol

UPW Ultra-Pure Water

RAT Reactor Analysis Tool (code package)

RATDB RAT Database

MC Monte-Carlo

IBD Inverse Beta Decay

ELLIE Embedded LED/Laser Light Injection Entity

AMELLIE ELLIE Attenuation Module

SMELLIE ELLIE Scattering Module

TELLIE ELLIE Timing Module

FV Fiducial Volume (the target detector volume)

^{214}BiPo ^{214}Bi and ^{214}Po coincident decays

AmBe ^{241}Am-^{9}Be source

PDF Probability Density Function

IAEA International Atomic Energy Agency

PWR Pressurised Water Reactor

PHWR Pressurised Heavy Water Reactor

LF Loading Factor

IESO Independent Electricity System Operator (Ontario)

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

BHT butylated hydroxytoluene

MSW Mikheyev-Wolfenstein-Smirnov

ROI Region Of Interest (for $0\nu\beta\beta$)

PR Proton Recoil