

Improving the Sensitivity to Neutrinoless Double Beta Decay in SNO+



Tereza Kroupová

Wolfson College

University of Oxford

A thesis submitted in fulfilment of the requirements

for the degree of

Doctor of Philosophy

Trinity Term 2020

Abstract

SNO+ is a multi-purpose neutrino detector located 2 km underground in Sudbury, Canada. The inside of the detector is currently being filled with liquid scintillator, which will then be loaded with tellurium-130 in order to search for neutrinoless double beta decay ($0\nu\beta\beta$).

This thesis concerns itself with ways to improve the sensitivity of SNO+ to this process and achieves it in two ways. Firstly, it is demonstrated that an increase of the light output of 15 – 20% can be achieved with the addition of an amine to the tellurium-loaded scintillator mixture. The higher light output results in better energy resolution of the detector, which conversely translates in higher sensitivity to $0\nu\beta\beta$. Furthermore, it is demonstrated that this addition is important for stabilisation of the tellurium-loaded scintillator with regards to water exposure. Secondly, the application of multi-site event classification techniques in signal extraction in SNO+ is discussed. A likelihood analysis including topological and timing pulse shape discrimination (PSD) parameters as additional observables is presented. These PSD parameters differentiate between interactions with point-like and multi-site energy depositions. Their use can therefore break the degeneracy between any "point-like" $0\nu\beta\beta$ signal and more distributed energy depositions from radioactive background decays involving γ s, such as from cosmogenic or external sources. The resulting expected lower limit on $0\nu\beta\beta$ half-life is 1.80×10^{26} years at 90 % confidence after 3 years of data taking. Furthermore, it is demonstrated that this analysis would allow for a genuine $0\nu\beta\beta$ observation for effective Majorana mass of 75 – 181 meV at 99 % confidence after 3 years, which is 40 – 50 % lower than what is achievable without the inclusion of PSD.

Acknowledgements

There are many people without whom this work would not have been possible. Firstly, I am very grateful to my supervisor, Steve Biller, for all his advice, guidance, encouragement and the opportunity to learn from him. His supervision has given me the freedom to become independent in my work, while he has always been there to provide feedback and new ideas.

Thank you Armin Reichold and Jeff Tseng for your additional support, inputs into my work and career, and the friendly academic environment you created. I would also like to thank Ed Leming, for all his time helping me with random issues, for teaching me some of his problem solving techniques and, importantly, for all the pep talks throughout the last 3.5 years.

I already miss sharing the SNO+ student office with the wonderful people that have become my friends: Josie Paton, Josh Wang, Iwan Morton-Blake, Jeff Lidgard, Esther Turner and Jack Dunger. Thank you for all the beer, quick (and less quick) work sanity checks, proof-reading and fun. Special thanks go to Jack as this thesis builds tremendously on his work and expertise, Esther for being such a wonderful friend and Iwan for his swearing that somehow makes every situation better.

I am lucky to have benefited from the large community of the sub-department of Particle Physics in Oxford with all its seminars and great infrastructure. I would especially like to thank the administrators Kim Proudfoot and Sue Geddes and the IT staff for all their help during my DPhil.

The SNO+ collaboration is an extremely friendly one and I have learned a lot from the its numerous members, for which I am thankful. It was always amazing to show up at SNOLAB, see so many familiar faces and discuss work. Thank you Steph and Leon, the happiest and grumpiest people in the world, respectively, for all the paddle-boarding and making my time in Sudbury an entirely different experience.

I would like to thank to my friends in Oxford, Prague and elsewhere, who all have been wonderful and made the past years a great fun. In particular, Eliška, Johanka and Vojta, who put up with all my complaints and always were there for me. Thank you Lucie for being an incredible housemate and for all the cake. Thank you Cam for being an awesome person and loving me even when we were locked down in one room together, trying to finish our theses.

Most importantly, I will always be grateful to my family for supporting me in whatever I decided to do. It meant an awful lot. My biggest thanks go to my mum Pavla, dad Ivan and brother Adam for all their love.

Contents

1	Introduction	1
2	Neutrinoless Double Beta Decay	3
2.1	The Neutrino	3
2.1.1	Evidence of Neutrino Mass	5
2.1.2	Size of Neutrino Mass and Mass Hierarchy	6
2.1.3	Neutrino Nature	7
2.1.4	The See-Saw Mechanism	9
2.2	Double Beta Decay	10
2.2.1	Neutrinoless Double Beta Decay Process	11
2.2.2	$0\nu\beta\beta$ Half-life	12
2.2.3	Experimental Signature of $0\nu\beta\beta$	13
2.2.4	Implications of $0\nu\beta\beta$	14
2.3	$0\nu\beta\beta$ Experiments	16
2.3.1	Semiconductors	17
2.3.2	Crystal Bolometers	19
2.3.3	Experiments with Tracking Detectors	20
2.3.4	Time Projection Chambers	21
2.3.5	Organic Liquid Scintillators	23
2.4	Choice of Double Beta Isotope	24
3	The SNO+ Experiment	28
3.1	Detector Overview	28
3.2	PMTs and Electronics	30
3.2.1	PMTs	30
3.2.1.1	Transit time	31
3.2.1.2	Wavelength efficiency	32
3.2.1.3	PMT Concentrators	33
3.2.2	Electronics	33

3.3	SNO+ Target and Phases	36
3.3.1	Water Phase	36
3.3.2	Unloaded Scintillator Phase	37
3.3.3	Tellurium Phase	39
3.4	Detector Calibration	39
3.4.1	Calibration Using Radioactive Sources	41
3.4.2	Optical Calibration	41
3.5	Software	44
4	Tellurium-Loaded Scintillator	45
4.1	Scintillation Process	45
4.1.1	Fluors	47
4.1.2	Light Yield	48
4.2	Scintillator in SNO+	49
4.3	Requirements on Tellurium Loaded Scintillator	51
4.4	Tellurium Loading	52
4.4.1	Tellurium Butanediol Synthesis	52
4.4.2	TeBD Loading Effect on Scintillator LY	53
4.4.3	TeBD Loading Effect on Scintillator Stability	57
4.5	Amine Addition	58
4.5.1	LY Improvement with DDA	60
4.5.2	Stability Improvement with DDA	61
4.5.2.1	Visual Inspection	61
4.5.2.2	LY Monitoring as a Measure of Stability	63
4.5.3	Mechanism	64
4.6	Underground Processing	65
4.7	Summary	66
5	Events in SNO+	67
5.1	Types of Events in SNO+	68
5.2	Event Reconstruction in Liquid Scintillator	71
5.2.1	Position and Time Reconstruction	71
5.2.2	Energy Reconstruction	73
5.3	Event Classification	77
5.3.1	Timing Classifier	78
5.3.2	Topological Classifier	85
5.4	Summary	89

6	Likelihood Analysis for Neutrinoless Double Beta Decay	92
6.1	Backgrounds for $0\nu\beta\beta$	93
6.1.1	$2\nu\beta\beta$	93
6.1.2	Solar Neutrino Scattering	94
6.1.3	Internal U/Th Natural Radioactivity	97
6.1.4	Cosmogenic Isotopes	101
6.1.5	External Backgrounds	103
6.1.6	AV Leaching and (α -n) Reactions	106
6.2	Analysis Framework	108
6.3	Monte Carlo Simulation and Fake Datasets	110
6.4	Observables	112
6.5	Analysis Cuts	114
6.6	Probability Density Functions	115
6.7	Verification of the Signal Extraction	119
6.8	Limit on $0\nu\beta\beta$ Half-life	121
6.9	Discovery Potential	124
6.10	Correlation with Cosmogenic Backgrounds	128
6.11	Summary	132
7	Conclusions and Outlook	134
A	List of Acronyms	137
B	List of Amines	140
C	^{60}Co and ^{208}Tl Decay Schemes	145
D	Detector State Correction	147
E	Tuning Classifier Weights on Different Backgrounds	152
F	Rates	156
	References	166

List of Figures

2.1	Neutrino mass hierarchies	7
2.2	Feynman diagrams for $2\nu\beta\beta$ and $0\nu\beta\beta$	12
2.3	Effective Majorana mass as a function of the lightest neutrino mass	14
2.4	Modelled energy spectra of $2\nu\beta\beta$ and $0\nu\beta\beta$	15
2.5	Double β decay Q-values as a function of natural abundance for selected double β isotopes	26
2.6	$2\nu\beta\beta$ half-lives for selected isotopes	27
3.1	Overview of the SNO+ detector	29
3.2	A schematic of the Hamamatsu R1408 PMT	31
3.3	Transit time probability distribution for a R1408 PMT	32
3.4	Wavelength efficiencies of R1408 PMTs	33
3.5	Drawing of a SNO PMT including housing and concentrator	34
3.6	A photo of water-scintillator interface inside the AV	38
4.1	π electrons in benzene	46
4.2	The chemical structures of LAB, PPO and bisMSB	49
4.3	Absorption and emission spectra of unloaded scintillator in SNO+	50
4.4	Synthesis of TeBD precursor molecules	53
4.5	Examples of possible compounds inside the TeBD complex	54
4.6	Comparison of charge spectra between unloaded and 0.5% Te- loaded scintillator	55
4.7	Photo of light yield monitoring setup	56
4.8	Block diagram of light yield monitoring setup	56
4.9	Light yield of LAB+PPO scintillator as a function of Te-loading concentration	58
4.10	Photos of an Te-loaded scintillator sample with and without flake formation	59
4.11	Chemical structure of N,N-dimethyldodecylamine	59

4.12	Light yield of Te-loaded scintillator as a function of DDA concentration	60
4.13	Photos of humidity exposed 1% Te scintillator samples over time	62
4.14	Effect of DDA addition on light yield stability after humidity exposure	64
5.1	Overview of event position reconstruction in liquid scintillator	74
5.2	Number of triggered PMTs as a function of electron energy	76
5.3	Overview of event energy reconstruction in liquid scintillator	77
5.4	Reconstructed energy of ^{60}Co , external ^{208}Tl events and $0\nu\beta\beta$ events inside the AV	79
5.5	Comparison of time residual spectra for ^{60}Co decay, e^- and $0\nu\beta\beta$ events	81
5.6	Time residual spectrum for external ^{208}Tl backgrounds	82
5.7	Fisher and $\Delta \log \mathcal{L}$ discriminants for $0\nu\beta\beta$, ^{60}Co and external ^{208}Tl events	84
5.8	Background acceptance as a function of signal efficiency for timing classifier cuts	86
5.9	A schematic of an external background event inside the SNO+ detector	87
5.10	Angular distribution of early PMT hits for signal and external background events	88
5.11	Topological discriminant values for signal and external ^{208}Tl	89
5.12	Background acceptance as a function of signal efficiency for topological classifier cuts	90
5.13	Contour plot of timing and topological discriminant values for $0\nu\beta\beta$ and external ^{208}Tl	91
6.1	Feynman diagrams for neutrino elastic scattering on electrons	95
6.2	Solar neutrino flux on Earth	95
6.3	The energy distribution of recoiled electrons from ^8B solar neutrino elastic scattering	96
6.4	The uranium-238 and thorium-232 decay chains	99
6.5	$^{212/214}\text{Bi}$ and $^{212/214}\text{Po}$ events before and after coincidence cuts	100
6.6	Reconstructed energy of cosmogenic isotopes produced from natural tellurium	104
6.7	Radial distributions of external ^{208}Tl and ^{214}Bi decays	105
6.8	Reconstructed energy of (α, n) interactions	107
6.9	Dependence of PSD discriminants on reconstructed energy	117

6.10	Dependence of PSD discriminants on reconstructed radial position	118
6.11	Fitted $0\nu\beta\beta$ counts for ensemble of sample datasets	122
6.12	Posterior probability of $0\nu\beta\beta$ counts for dataset with no signal . .	123
6.13	Posterior probability of $0\nu\beta\beta$ counts for Asimov datasets with signal	125
6.14	Fit result for an Asimov dataset with $0\nu\beta\beta$ signal	126
6.15	Comparison of fit results from signal extraction with and without PSD	127
6.16	Discovery potential for $0\nu\beta\beta$ in SNO+	127
6.17	Correlation of $0\nu\beta\beta$ posterior probability with ^{22}Na , ^{60}Co and ^{88}Y	129
6.18	Half-life sensitivity with increased ^{60}Co background	131
6.19	Posterior probability of $0\nu\beta\beta$ counts for Asimov datasets with ^{60}Co contamination and no signal	132
B.1	Photo of precipitation in a sample with tetradecylamine	141
C.1	^{60}Co decay scheme	145
C.2	^{208}Tl decay scheme	146
D.1	Reconstructed energy bias with a fraction of PMTs turned off at random	150
D.2	Reconstructed energy bias with PMT crates turned off	151
E.1	Performance of the timing classifier tuned on different backgrounds	153
E.2	Performance of the topological classifier tuned on different back- grounds	154

List of Tables

2.1	Current best limits on $0\nu\beta\beta$ half-life	18
3.1	Radioactive calibration sources in SNO+	42
6.1	Expected internal U/Th background rates in tellurium phase	97
6.2	Potential cosmogenic backgrounds to $0\nu\beta\beta$ search in SNO+	102
6.3	Number of $0\nu\beta\beta$ events in Asimov datasets	112
6.4	Cuts used in $0\nu\beta\beta$ likelihood analysis	114
6.5	Signal extraction bias for different versions of the fit	120
6.6	Signal extraction pull for different versions of the fit	121
6.7	Limits of $T_{1/2}^{0\nu\beta\beta}$ at 90% C.I.	123
6.8	Datasets for studies with increased ^{60}Co background	130
F.1	Fit parameters and expected background rates	156

Chapter 1

Introduction

Neutrinos are one of the least understood of the known elementary particles and could provide an insight into questions such as dark matter or the ratio of matter and anti-matter in the Universe. The discovery of neutrino oscillations by the Sudbury Neutrino Observatory and SuperKamiokande experiments [1,2] confirmed that neutrinos are massive particles and gave the only direct evidence of physics beyond the Standard Model to date. Since then, vast progress has been made in the measurement of the neutrino mixing angles and square mass differences. Despite that, our knowledge about neutrinos is still limited. The absolute mass scale, the neutrino mass ordering, and the role that neutrinos played in the early Universe remain unknown.

One of the biggest questions in fundamental particle physics is whether neutrinos are Dirac fermions, with distinct anti-particles, or Majorana fermions, for which the particles and anti-particles are identical. The best available probe of the neutrino nature is neutrinoless double beta decay ($0\nu\beta\beta$), a hypothetical process that require massive Majorana neutrinos. The discovery of $0\nu\beta\beta$ would not only confirm the Majorana nature of neutrinos, but also provide a measurement of the absolute neutrino mass scale and bolster various theories that seek to go beyond the Standard Model [3–5].

The focus of this thesis is the study of $0\nu\beta\beta$ sensitivity in the SNO+ experiment, which will use natural tellurium for this search. In particular, methods of improving the sensitivity to this process through both detector target R&D and signal extraction developments are explored.

A brief introduction to neutrino physics, the process of neutrinoless double beta decay and the experimental searches that seek to detect this process is given in Chapter 2. The SNO+ experiment is introduced in Chapter 3.

Chapter 4 focuses specifically on the SNO+'s liquid scintillator mixture into which tellurium is incorporated and the improvements of its stability and light output that are achieved with the addition of an amine. Chapter 5 then describes particle interactions inside the scintillator volume, their detection and reconstruction. Classification of interactions based on their timing and topological distributions is also described there.

A likelihood analysis for the neutrinoless double beta search is presented in Chapter 6. Particular attention is paid to the improvements achieved by including the classification parameters from previous chapter as additional observables in the signal extraction.

Finally, the findings from this work are summarised in Chapter 7.

Chapter 2

Neutrinoless Double Beta Decay

The existence of neutrinoless double beta decay ($0\nu\beta\beta$) is presently one of the main subjects of investigation in particle physics with a number of scientific collaborations devoting their attention to the search for this process. A brief introduction of neutrinos is provided here, together with the open questions surrounding their mass, which motivate the search for $0\nu\beta\beta$. This hypothesised process is then described, including the expected experimental signature and decay half-life. The later sections provide an overview of the current state of $0\nu\beta\beta$ research and considerations for experiment design.

2.1 The Neutrino

Neutrinos are elementary particles and their existence was confirmed in 1956 by F. Reines and C. Cowan [6]. They form a part of the Standard Model of Particle Physics (SM), where they belong to the category of leptons - spin $\frac{1}{2}$ particles that do not participate in the strong interaction. There are three families of leptons in the SM: e , μ and τ , each containing a charged lepton (e^- , μ^- , τ^-) and an

associated neutrino (ν_e, ν_μ, ν_τ). The family to which a lepton belongs is said to determine its "flavour".

While charged leptons participate in electromagnetic interactions, neutrinos have no charge and only interact weakly. The weak force is mediated by massive gauge bosons, namely W^+ , W^- and Z . It is the high mass of these bosons that lead to a steep decrease in the strength of this force over distance. Consequently, the cross-section of neutrinos with matter is very small ($\mathcal{O}(10^{-38} \text{ cm}^2 \text{ per nucleon per GeV of neutrino energy})$ [7]), which makes them particularly difficult to detect.

On the other hand, neutrinos are abundant in the Universe and are produced in a variety of processes, such as: thermonuclear reactions in stars, fission in nuclear reactors, decays of naturally occurring radioactive isotopes, supernova explosions and interactions of cosmic rays in the Earth's atmosphere. In these interactions, neutrinos are created in a 3-prong vertex between the neutrino, a weak gauge boson and a charged lepton. The flavour of the neutrino is determined by the charged lepton as the weak interaction conserves lepton flavour in the SM. A lepton number, L , is defined to have a value of 1 for leptons and -1 for anti-leptons. It is also thought to be conserved in the SM and no interactions violating this number have been observed to date. However, there is no fundamental reason for this symmetry and various new theories propose its violation.

In the SM, neutrinos are massless. Furthermore, the weak force only interacts with particles of left-handed chirality and anti-particles of right-handed chirality. Helicity and chirality are equivalent for relativistic particles and, therefore, only neutrinos of left-handed helicity and right-handed anti-neutrinos exist in the model [8].

2.1.1 Evidence of Neutrino Mass

The discovery of neutrino oscillations [1,2] proved that neutrinos have mass and, thus, provided the first laboratory observation of physics Beyond the Standard Model (BSM). In this process, neutrinos change their flavour as they propagate and, therefore, can be detected with a flavour different to the one in which they were created. This is possible because the neutrino flavour eigenstates $|\nu_\alpha\rangle$ ($\alpha = e, \mu, \tau$) are a superposition of mass eigenstates $|\nu_i\rangle$ ($i = 1, 2, 3$). The relation between these eigenstates is described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, U :

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i} |\nu_i\rangle \quad U = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix}. \quad (2.1)$$

If the mass eigenvalues are distinct from each other, the mass eigenstates propagate through space with different phases. As a result, the superposition of these states changes over time and, therefore, the neutrino can be detected as a different flavour. The probability of such event is the described by the oscillation probability, $P_{\alpha \rightarrow \beta}$, approximated as:

$$P_{\alpha \rightarrow \beta}(E, L) = \sum_{i,j} U_{\alpha i}^* U_{\beta i} U_{\alpha j} U_{\beta j}^* \exp\left(-i \frac{\Delta m_{ij}^2 L}{2E}\right), \quad (2.2)$$

where L is the distance between the place where the neutrino was created and where it is detected. E is the energy of the neutrino and $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$ is the squared-mass difference between eigenstates ν_i and ν_j . The derivation of equation 2.2 can be found in [8].

2.1.2 Size of Neutrino Mass and Mass Hierarchy

Many of the components of the PMNS matrix have now been precisely measured in various neutrino oscillation experiments [9]. Importantly, the existence of neutrino oscillations means that at least two out of the three neutrinos are massive and the squared-mass differences have also been measured.

By convention, Δm_{21}^2 is defined to be the smaller of the squared-mass differences. With ν_1 set to have the largest contribution to ν_e , the difference has been measured to be positive with $m_2 > m_1$ [10]. The measurement of this quantity comes from solar neutrino and reactor anti-neutrino measurements and is sometimes referred to as Δm_{sol}^2 . The other independent squared-mass difference, Δm_{32}^2 , which is an order of magnitude larger than Δm_{21}^2 , is measured mostly from atmospheric neutrinos and, hence, it is denoted as Δm_{atm}^2 . Unfortunately, only the absolute value of this difference, $|\Delta m_{atm}^2|$, can be measured and, therefore, there are two possible neutrino mass orderings: $m_1 < m_2 \ll m_3$ referred to as "normal hierarchy" (NH) and $m_3 \ll m_1 < m_2$ referred to as the "inverted hierarchy" (IH). The two scenarios are also graphically depicted in Figure 2.1. It is noted, that strictly speaking, $|\Delta m_{atm}^2|$ corresponds to $|\Delta m_{32}^2|$ in the NH case and to $|\Delta m_{31}^2|$ in the IH case, however the uncertainties on these measurements are much larger than Δm_{21}^2 and, therefore, $|\Delta m_{32}^2| \approx |\Delta m_{31}^2|$.

Despite a vast amount of data collected by oscillation experiments, the absolute values of the neutrino masses remain unknown as these experiments are not sensitive to it ¹. Upper limits on the mass of the electron neutrino come from direct neutrino mass experiments undertaking kinematic studies of weak-processes, but the precision of these measurements is not high enough to give the absolute value for the mass. The current best limit comes the measurement of β decay energy spectrum endpoint of ${}^3\text{H}$ and constrains the mass to < 1.1 eV at 90 % confidence level [12].

¹apart from a lower limit, which can be obtained by setting the lowest mass eigenvalue to 0

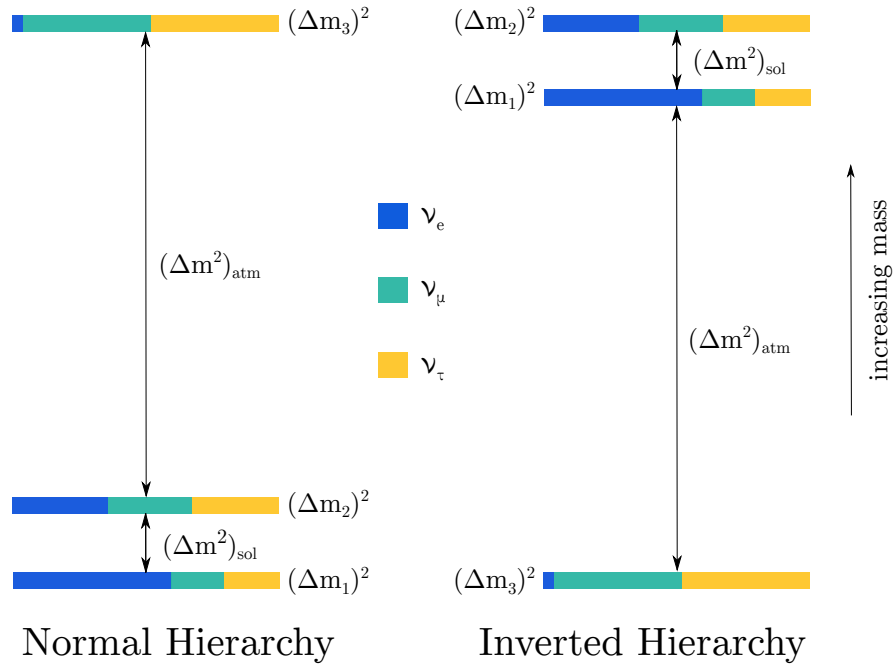


Figure 2.1: The two possible neutrino mass orderings, with NH on the left and IH on the right. Mass is increasing in the direction towards the top edge of the figure and the different colours depict the favour content of each mass eigenstate taken from the current best measurements of the PMNS matrix components. Figure modified from [11].

Finally, as neutrinos from the Big Bang played a crucial role in the formation of large scale structures in the Universe, constraints on neutrino masses can, in principle, be obtained from cosmological observations. In this case, the sum of the neutrino masses is measured and a recent result constrained this quantity to < 1.2 eV [13]. Tighter constraints are expected in the future. However, the results of cosmological surveys are highly dependent on the cosmological model used in the data analysis.

2.1.3 Neutrino Nature

Another puzzle surrounding the neutrino mass is its origin. In the SM model, fermions gain mass through the Higgs mechanism [8]. Here, mass is generated by the Higgs field interacting with both right-handed and left-handed chiral

components of the particle. The corresponding term in the Lagrangian is

$$\mathcal{L}_{mass}^D = -\frac{gv}{\sqrt{2}}(\bar{\psi}_R\psi_L + \bar{\psi}_L\psi_R), \quad (2.3)$$

where ψ is the given particle field, subscripts denote the left and right chiral components of that field and bar denotes an adjoint of the field component. g is a constant describing the coupling of the particle field to the Higgs field and v is the vacuum expectation value of the Higgs field. The resulting mass of the particle is given by

$$m_D = \frac{gv}{\sqrt{2}}. \quad (2.4)$$

Particle species that gain their mass through this mechanism and are distinct from their anti-particles are called Dirac particles. It is clear from equation 2.3 that in order for a neutrino to be Dirac and gain mass in the same way as the other fermions, both right and left-handed chiral states must exist. As only left-handed neutrinos couple to the weak force, a right-handed "sterile" neutrino is introduced, which only interacts through gravity and, hence, has never been observed. This addition to the SM, treating neutrinos as massive Dirac particles, is known as the "minimal extension of the SM".

A problematic aspect of giving neutrinos mass as described above is its small size. The neutrino is at least 5 orders of magnitude lighter than the next lightest fermion, which means that the Higgs coupling to neutrinos would have to be extremely small compared to the couplings for other particles.

An alternative mechanism was proposed in 1937 by E. Majorana [14], who suggested that for some particles

$$\psi^C = \psi, \quad (2.5)$$

where $\psi^C = i\gamma^2\psi^*$ is the charge conjugate of ψ and γ^2 is one of the four-dimensional Dirac gamma matrices. As the charges of ψ and ψ^* have opposite

magnitudes, the above equation can only hold for neutral particles. Neutrinos are the only known neutral fermions and, therefore, a good candidate to obey Majorana's proposed equality. If the condition from equation 2.5 is applied, the right-handed neutrinos appear as charge-parity conjugates of left-handed anti-neutrinos. A Majorana mass term can be added to the SM Lagrangian:

$$\mathcal{L}_{mass}^L = -m_L \left(\bar{\psi}_L^C \psi_L + \bar{\psi}_L \psi_L^C \right), \quad (2.6)$$

where m_L is the Majorana mass of left-handed neutrino. This term provides direct coupling between the neutrino and anti-neutrino fields, which can be interpreted as neutrinos being their own anti-particles. In this case, neutrinos are called Majorana particles as opposed to Dirac particles.

2.1.4 The See-Saw Mechanism

The existence of neutrino Majorana mass term, does not prohibit the presence of Dirac mass and right-handed neutrino fields. In fact, the most general Lagrangian contains both of these contributions. This way, the neutrino would gain mass by both coupling to the Higgs field and through the mixing of the particle and anti-particle fields. This again requires right-handed neutrino fields and, therefore, the right-handed equivalent of equation 2.6 is also present in the general neutrino mass Lagrangian. The right-handed neutrinos emerging from this combined Lagrangian are, however, different from the sterile light neutrinos necessary to generate neutrino mass purely by the Higgs mechanism. One of the simplest theories assumes

$$m_D \ll m_R \quad m_L = 0, \quad (2.7)$$

which gives rise to two particles: a light left-handed neutrino (ν) and a heavy right handed neutrino (N) with their respective masses

$$m_\nu \simeq \frac{m_D^2}{m_R} \qquad m_N \simeq m_R. \qquad (2.8)$$

This procedure is known as the "see-saw mechanism"² and a more in-depth description of it can again be found in [8]. Importantly, the small neutrino mass is explained by the large size of m_R , which is related to high-energy scale of new physics beyond the SM. The Dirac mass, m_D , can be of the same order as for the same fermions without the need for "unnaturally" small coupling. Therefore, the see-saw mechanism provides a compelling explanation for the neutrino mass scale. In order to confirm it, however, the Majorana nature of neutrinos must first be experimentally established.

The most promising strategy for resolving whether neutrinos are Dirac or Majorana particles would be the observation of neutrinoless double beta decay [16]. This process and the implications of its discovery are described in the following section.

2.2 Double Beta Decay

In some nuclear isotopes, with a proton number Z and a nucleon number A , a single β decay

$$(Z, A) \rightarrow (Z + 1, A) + e^- + \bar{\nu}_e \qquad (2.9)$$

is energetically unfavourable because the combined mass of the daughter nucleus and the electron are larger than the mass of the parent. A double β decay

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- + 2\bar{\nu}_e \qquad (2.10)$$

²other versions of this mechanism with $m_L \neq 0$ are also discussed in literature [15]

might, however, be allowed. In this process, two neutrons from the same nucleus simultaneously undergo a β decay, each resulting in the production of a proton, an electron and an anti-neutrino. This decay is referred to as two neutrino double β decay ($2\nu\beta\beta$). It was first proposed by M. Goeppert-Mayer [17] in 1935 and has now been observed in a number of isotopes where single β decay is forbidden or strongly suppressed [7].

A typical double β candidate is a nucleus with even number of both protons and neutrons. Such even-even nuclei have strong bonding between the paired nucleons, which is responsible for a comparatively lower mass of these nuclei. The daughter nucleus that would be created by a single β transition is an odd-odd nucleus and, hence, it is less likely to have a lower mass than the parent due to this pairing energy. On the other hand, a double β decay changes the nucleon number by two units and again creates an even-even nucleus. Some of the isotopes in which $2\nu\beta\beta$ was observed are discussed further in section 2.4.

Double β decay is a second order weak process and, therefore, the observed half-lives are long ($\mathcal{O}(10^{18} - 10^{21}$ years)) [7]. The total kinetic energy released during the transition from the parent nuclei ground state to the daughter ground state is the Q-value, $Q_{\beta\beta}$. In a $2\nu\beta\beta$, part of this energy is carried away by the two anti-neutrinos, which are not detected in an experiment. Therefore, the visible energy of this process is a continuous spectrum with an endpoint at the $Q_{\beta\beta}$.

2.2.1 Neutrinoless Double Beta Decay Process

Unlike $2\nu\beta\beta$, neutrinoless double β decay ($0\nu\beta\beta$),

$$(Z, A) \rightarrow (Z + 2, A) + 2e^- \quad (2.11)$$

is a hypothetical process that has never been observed [16]. While $2\nu\beta\beta$ is allowed in the Standard Model, $0\nu\beta\beta$ violates lepton number conservation by two

units and requires neutrinos to be Majorana particles. In this case, there are no anti-neutrinos emitted from the nucleus during the decay as shown in Figure 2.2³. Instead, the process is mediated by a virtual Majorana neutrino. It can be thought of as an anti-neutrino emitted at one vertex and a neutrino absorbed at the other, which is only possible if the particle and anti-particle are not distinct, as is the case for Majorana particles. Furthermore, in order to conserve momentum, the virtual neutrino must flip its helicity between the two vertices, which is only achievable for particles with non-zero mass.




Figure 2.2: Feynman diagrams for $2\nu\beta\beta$ and $0\nu\beta\beta$. The $2\nu\beta\beta$ process is shown on the left and results in the production of a daughter nucleus (with proton number increased by 2), 2 electrons and 2 anti-neutrinos. On the right, $0\nu\beta\beta$, where no neutrinos are produced, is depicted. The latter is only allowed if neutrinos are Majorana particles.

2.2.2 $0\nu\beta\beta$ Half-life

The half-life of $0\nu\beta\beta$, $T_{1/2}^{0\nu}$, is directly related to the effective Majorana mass, $\langle m_{\beta\beta} \rangle$, by

$$\left(T_{1/2}^{0\nu}\right)^{-1} = \langle m_{\beta\beta} \rangle^2 \times |M_{0\nu}|^2 \times G_{0\nu}, \quad (2.12)$$

³There is a number of possible $0\nu\beta\beta$ mechanisms, which can also involve yet undiscovered particles [18, 19]. In this work, only the light neutrino exchange, which is shown in the figure, is discussed.

where $M_{0\nu}$ is the nuclear matrix element of the given isotope and $G_{0\nu}$ is the phase-space factor [20]. The effective Majorana mass is defined as

$$\langle m_{\beta\beta} \rangle = \left| \sum_i m_i U_{ei}^2 \right|, \quad (2.13)$$

with i being the number of neutrino mass eigenstates, m_i the mass eigenvalues and U is the PMNS matrix transforming between mass and flavour eigenstates from equation 2.1. The sum over the matrix elements comes from the fact that the electron (anti-)neutrino is a superposition of mass eigenstates and, hence, the probability of a W boson in the $0\nu\beta\beta$ to interact with a mass m_i is proportional to U_{ei}^2 . The factor of m_i in equation 2.13 arises from the aforementioned helicity flip, which is less likely to happen for lower masses. It is due to this helicity suppression that the half-life of $0\nu\beta\beta$ is expected to be even longer than in the two neutrino case, which makes the detection of this process very difficult. Indeed, the current best lower limit on $T_{1/2}^{0\nu}$ is 1.07×10^{26} years at 90 % confidence, corresponding to $\langle m_{\beta\beta} \rangle = 61 - 165$ meV [21]. Figure 2.3 shows the effective Majorana mass as a function of the lightest neutrino mass eigenstate together with this limit. A large portion of the allowed values of $\langle m_{\beta\beta} \rangle$ remains unprobed.

$M_{0\nu}$ and $G_{0\nu}$ in equation 2.12 are both nuclear factors. Notably, $G_{0\nu}$ depends on both Z and the Q -value, which is important when choosing a suitable isotope for a $0\nu\beta\beta$ search, as is further discussed in section 2.4. It is crucial to state that the nuclear matrix elements are a source of large uncertainty and their theoretical values depend on the nuclear model used. Therefore, when converting between $T_{1/2}^{0\nu}$ and $\langle m_{\beta\beta} \rangle$, the used value of the matrix element must always be stated.

2.2.3 Experimental Signature of $0\nu\beta\beta$

In contrast to the two neutrino case, the entire energy that is released in $0\nu\beta\beta$ is detected, because no portion of it is carried away by neutrinos. As a result,




Figure 2.3: Effective Majorana mass as a function of the lightest neutrino mass in the case of 3 active neutrinos. Blue and green regions show the allowed values in the case of inverted and normal hierarchy, respectively. The current best upper limit on $\langle m_{\beta\beta} \rangle$ from [21] is shown in orange. Figure modified from [22].

the signature of this process is a peak in the energy spectrum centred at the Q -value as illustrated in Figure 2.4. In principle, the peak is a delta function, but in any realistic experiment it will be smeared by the detector resolution as also demonstrated in Figure 2.4. In addition to the overall observed energy, $0\nu\beta\beta$ can also be distinguished by two electrons emerging from the same vertex. The strategies employed by various experiments to detect this process are discussed further in section 2.3.

2.2.4 Implications of $0\nu\beta\beta$

If $0\nu\beta\beta$ is discovered, the Majorana nature of neutrinos, required by many BSM theories, will be confirmed. The neutrino mass generation will be explainable without the need of fine tuning of the Higgs coupling to neutrinos. The effective Majorana mass will also be measured. Even though this cannot be directly translated to the absolute masses of neutrinos, it might be possible to distinguish




Figure 2.4: Modelled energy spectra of $2\nu\beta\beta$ (dashed line) and $0\nu\beta\beta$ (full line). The rate of $0\nu\beta\beta$ signal is simulated as 1 % of $2\nu\beta\beta$ rate. In reality, the signal is expected to be smaller. Figure taken from [20].

between the normal and inverted hierarchies depending on the value. Additionally, this process would provide the first evidence of lepton number violation.

Finally, the discovery would hint at the existence of heavy Majorana neutrinos, which could explain the predominance of matter in our Universe through process called leptogenesis [3]. It is assumed that equal amounts of matter and anti-matter were created in the Big Bang. However, with the current understanding of physics, all the matter and anti-matter would annihilate to create photons. Cosmological surveys have established that only a small excess of quarks over anti-quarks (1 part in 10^9 [23]) must have been present in the early Universe. Nevertheless it allowed for the formation of all the stars, galaxies and other cosmic structures. How this asymmetry was created is one of the biggest questions in modern physics. If neutrinos are Majorana particles, two Majorana phases will be present in the PMNS matrix, which can be a source of charge-parity (CP) violation. If these phases are non-zero, the heavy right-handed neutrinos can decay preferentially to charged leptons rather than anti-leptons. This would create an asymmetry between leptons and anti-leptons that can be then transferred into the baryon sector [24].

All these potential implications of $0\nu\beta\beta$ provide a compelling reason to look for this process. The rest of this chapter is dedicated to the experimental techniques used for this search.

2.3 $0\nu\beta\beta$ Experiments

The experiments striving to discover $0\nu\beta\beta$ search for a rare peak at the expected Q-value on top of a continuous energy spectrum. This background spectrum comes from variety of sources: the tail of $2\nu\beta\beta$, natural radioactivity from inside and outside of the detector as well as any other interactions that can be mistaken for $0\nu\beta\beta$ signal. Therefore, the most important consideration in designing a $0\nu\beta\beta$ experiment is optimising the signal with respect to the background.

Assuming Poisson distribution of the background, negligible systematic uncertainty, and a set energy range, the uncertainty on the number of background counts, B , is \sqrt{B} . The sensitivity of a background limited experiment can then be estimated in terms of the number of standard deviations as $c = S/\sqrt{B}$, where S is the number of signal counts. Given a number of double β nuclei $N_{\beta\beta}$, signal detection efficiency ϵ , experiment runtime t and $0\nu\beta\beta$ half-life $T_{1/2}$, the expected number of observed signal counts is

$$S = \epsilon N_{\beta\beta} \frac{\ln(2)}{T_{1/2}} t. \quad (2.14)$$

Furthermore, with a constant background rate b , $B = bt$. The expected half-life sensitivity at confidence c is therefore

$$T_{1/2}^c = \frac{\ln(2)}{c\sqrt{b}} \epsilon N_{\beta\beta} \sqrt{t}. \quad (2.15)$$

Hence, to achieve significant sensitivity, the following is needed: large amount of isotope, good detection efficiency, long data taking time and low background

rates. Moreover, the background rate depends on the considered energy region of interest (ROI), with wider regions resulting in more background counts. The width of the $0\nu\beta\beta$ peak and, hence, the ROI depends on the energy resolution. The "spill-over" of the steep edge of $2\nu\beta\beta$ spectrum into the ROI also increases as resolution worsens. As a result, the increase in b scales non-linearly with energy resolution. Therefore, the achievable energy resolution is also a crucial factor in designing a $0\nu\beta\beta$ experiment.

In general, $0\nu\beta\beta$ experiments are conducted deep underground to avoid cosmic ray backgrounds and strive to use materials with extremely low radioactive contamination. Smaller experiments tend to have lower background rates and better energy resolution. Spatial resolution is also typically better, which can help in background rejection. However, smaller detectors typically contain less double β isotope. On the other hand, it is difficult to maintain good resolution in large detectors and scaling up experiments made up from smaller modules is often expensive. Cost, safety and technical feasibility also play important roles in experiment design.

Various scientific collaborations take different experimental approaches to achieve the best sensitivity. The current best lower limits on $0\nu\beta\beta$ half-life are listed in Table 2.1⁴. The existing and planned experiments searching for $0\nu\beta\beta$ can be classified either by the double β isotope or the experimental design. The latter approach is employed in the non-exhaustive summary of the field below.

2.3.1 Semiconductors

Semiconductor detectors are an established and sensitive technique for the detection of ionising radiation. The small energy band gap between valence and conduction electrons in semiconductors provides a good signal-to-noise ratio upon particle passage and an excellent energy resolution. Germanium semi-

⁴However, it has to be noted that these limits translate to different values of $\langle m_{\beta\beta} \rangle$ due to the different nuclear factors for the various isotopes, further discussed in section 2.4

Isotope	$T_{1/2}^{0\nu}$ Limit [10^{25} y]	Experiment	Isotope Exposure [kg·y]	Isotope Mass [kg]
^{48}Ca	5.8×10^{-3} [25]	ELEGANT VI	0.36	0.175
^{76}Ge	9 [26]	GERDA	82.4	35.6
^{82}Se	3.6×10^{-2} [27]	NEMO-3	3.5	0.9
^{96}Zr	9.2×10^{-4} [27]	NEMO-3	0.035	0.009
^{100}Mo	0.11 [28]	NEMO-3	34.7	6.9
^{116}Cd	1.0×10^{-2} [29]	NEMO-3	2.16	0.410
^{130}Te	3.2 [30]	CUORE	372.5	206
^{136}Xe	10.7 [21]	KamLAND-Zen	504	344
^{150}Nd	2.0×10^{-3} [31]	NEMO-3	0.19	0.037

Table 2.1: Current best limits on $0\nu\beta\beta$ half-life for different isotopes. All results are quoted at 90% confidence. Last two columns give the isotope exposure and mass with which these limits were measured, respectively.

conductor detectors are particularly useful for $0\nu\beta\beta$ searches, because ^{76}Ge is a double β isotope. Therefore, germanium can serve as a detector and source material simultaneously, which results in high signal efficiency. Taking advantage of the well-tested technology, the Heidelberg-Moscow and IGEX experiments were the first ones to achieve sensitivity to the $0\nu\beta\beta$ half-life of $\mathcal{O}(10^{25}$ y) [32, 33].

The currently operating experiment with the best limit on $T_{1/2}^{0\nu}$ in ^{76}Ge , GERDA [34], uses germanium detectors enriched in this isotope to 86 – 88%. The detector array has a total mass of 35.6 kg, and is immersed in liquid argon to provide passive shielding from ionising radiation originating outside of the detectors [35]. Furthermore, liquid argon scintillates and, hence, can be used to actively veto backgrounds that deposit energy inside this shield. The GERDA collaboration also employs a digital analysis of the detector signal pulses to discriminate between single-site energy deposition, such as the one expected from the $0\nu\beta\beta$ daughter electrons, and more distributed deposition from background

events involving γ s. The combination of high radiopurity of the detector material, good shielding, active background rejection through liquid argon and pulse shape discrimination results in an exceptionally low background count in the ROI.

Another currently running experiment, the Majorana Demonstrator [36], operates germanium detectors enclosed in two copper vacuum cryostats. The collaboration was able to achieve energy resolution of $\sim 0.05\%$ at $Q_{\beta\beta}$, the best resolution achieved in any $0\nu\beta\beta$ experiment so far.

The drawbacks of using semiconductors are the high cost of detector fabrication and difficult scalability inside the cryostat. Furthermore, the technology dictates the choice of the double β isotope. Nevertheless, a next generation experiment, LEGEND [37], will combine developments by the GERDA and Majorana collaborations and is expected to use one tonne of germanium detectors.

There are also ongoing efforts in alternative semiconductor technologies for $0\nu\beta\beta$ studies with other isotopes, such as CdZnTe diodes [38]. However, in terms of competition with other leading experiments, they are currently in early stages of development.

2.3.2 Crystal Bolometers

Bolometric detectors operate at very low temperatures (~ 10 mK) and use calorimetric approach for measuring energies of nuclear interactions [39]. The temperature of the absorber material containing the double β isotope is constantly monitored by a thermometer. When a nuclear decay happens inside the absorber, its temperature rises as the daughter particles deposit energy in the material. Crystalline materials have particularly low heat capacity at cryogenic temperatures, resulting in measurable temperature changes induced by small energy deposits (~ 0.1 mK/MeV) [39].

Contrary to semiconductor technology, it is in principle possible to use a

variety of crystals as absorbers and search for $0\nu\beta\beta$ in a number of different isotopes. To date, however, bolometers have mainly been used to search of $0\nu\beta\beta$ in ^{130}Te by the CUORE experiment and its prototypes [40]. CUORE uses cubic TeO_2 absorber crystals with 5 cm sides. With 988 such crystals enclosed inside a multi-layer cryostat cooled with helium, the total mass of the double β isotope is 206 kg.

CUORE has demonstrated the excellent energy resolution of this technique of $\sim 0.1\%$ at $Q_{\beta\beta}$ [41], which is comparable to germanium diodes. As it is possible to increase the isotope mass by simply adding more crystal units, the scalability is mostly limited by the size of the cryostat. The main drawback of crystal bolometers is the lack of reliable particle identification (i.e. distinguishing between α , β and γ particles), which leads to poor background rejection. The absorber crystals can be grown with high radiopurity. However, radioactive decays from the ^{238}U and ^{232}Th chain contamination on the surface of the crystals and other parts inside the cryostat are problematic and limit the sensitivity of these experiments.

The planned next generation experiment, CUPID [42], proposes to address this issue with scintillating crystals. In this approach, passage of charged particles results in photon emissions in addition to the aforementioned temperature increase. The shape of the recorded light pulse can be used for particle identification.

2.3.3 Experiments with Tracking Detectors

One approach to reject backgrounds is to reconstruct the two daughter electrons from the double β decay emerging from a single vertex. This strategy is employed by the NEMO-3 experiment [43] and its successor SuperNEMO [44], both of which contain large tracking detectors to record the trajectories of charged particles. The trackers are wire drift chambers filled with a he-

lium/argon/ethanol gas mixture that is ionised upon particle passage. The double β isotope is localised on a source foil and sandwiched between the tracking detectors, which are in turn surrounded by plastic scintillator calorimeter blocks to measure the energy of the electrons. As a result, the NEMO-3 and SuperNEMO possess modest energy resolution and excellent background rejection through vertex reconstruction.

Another advantage of the modularity of these detectors is the possibility of exchanging the double β isotope by simply replacing the source foil. This was indeed exploited in NEMO-3, which consisted of 20 modules taking data with 7 different isotopes. Despite no observed $0\nu\beta\beta$ signal, it still holds the best limit on the $0\nu\beta\beta$ half-life for 5 out of these 7 isotopes. Importantly, it was the first experiment to obtain precise measurements of $2\nu\beta\beta$ half-lives for many of these isotopes [45–47,31], which provided useful input for other experiments.

The disadvantage of this technology is that the source foil has to be as thin as possible ($\sim 0.02 - 0.3$ mm depending on the isotope [48]) to avoid scattering of the electrons inside it. Therefore, large area is needed to achieve high isotope exposure making this technology expensive compared to others. The SuperNEMO "demonstrator module" has dimensions $4 \times 6 \times 2$ m and many copies of this detector would be needed to reach the sensitivity of other next-generation $0\nu\beta\beta$ experiments. However, should a $0\nu\beta\beta$ signal be detected in another experiment, SuperNEMO can use its unique vertex reconstruction to unambiguously confirm the discovery. Furthermore, SuperNEMO would be able to shed light on the $0\nu\beta\beta$ mechanism by measuring the kinematic distributions of the outgoing electrons.

2.3.4 Time Projection Chambers

A time projection chamber (TPC) consists of a detector volume placed inside an electric field. Charged particles ionise the detector material and the resulting

electrons drift in the electric field towards a read-out system, typically consisting of anode multi-wire planes [49]. The total number of electrons collected on the anode is related to the energy of the particle and the locations of signal collection on the wire planes give a 2D image of the particle trajectory.

In some cases, the material filling the TPC can also be selected to have scintillating properties (i.e. emit light upon the passage of ionisation radiation) and the resulting photons can then be detected using light detectors. The scintillation signal appears before the anode pulse, which is delayed proportionally to the distance the ionisation electrons have to drift to reach the anode. The timing between the two signals supplements the 2D wire plane image to create a 3D reconstruction of the particle passage. Additionally, the relative amplitudes of the two signal pulses provides information about the type of the particle in the detector, with low ionisation to scintillation ratio for α particles and high ionisation to scintillation for β and γ particles. The ability to distinguish between particle types using the signal ratios as well as spatial imaging leads to excellent background rejection capabilities of TPCs. The energy resolution of the current TPC technology is worse compared with semiconductors and bolometers, but is compensated by better scalability.

For $0\nu\beta\beta$ searches, xenon enriched in the double β isotope ^{136}Xe is a convenient choice of detector material that also serves as a source. Xenon is an inorganic scintillator. Furthermore, it is chemically inert, which ensures that the ionisation electrons can drift sufficient distances in the electric field to reach the anode. Good results have been achieved by the EXO-200 experiment [50] using a liquid xenon TPC. EXO-200 set a $T_{1/2}^{0\nu}$ limit of 1.8×10^{25} years with $\mathcal{O}(100 \text{ kg})$ of xenon [51]. The next generation experiment, nEXO [52], is planning to deploy 5 tonnes of this element in a single TPC.

A slightly different approach is taken by another future $0\nu\beta\beta$ experiment, NEXT [52], which aims to build a TPC filled with high-pressure xenon gas.

Compared with liquid xenon, electron trajectories are longer in xenon gas, which might enable identification of the two electrons from a double β decay.

Additionally, both the nEXO and NEXT collaborations are investing significant effort into developing a technique for tagging the ^{136}Ba nuclei that are left over after the double β decay of ^{136}Xe [53, 54]. Identification of barium would result in elimination of all backgrounds apart from the tail of $2\nu\beta\beta$.

2.3.5 Organic Liquid Scintillators

Another strategy for obtaining good sensitivity to $0\nu\beta\beta$ is to maximise the exposure using large homogeneous detectors with hundreds of kilograms of isotope. Organic liquid scintillators are a convenient choice, because they can be purified to a high degree to reduce radioactive background contaminants [55]. Furthermore, the purification can be performed even during data taking by recirculating the organic scintillator through a purification system. Similarly, the circulation can be used to gradually mix ("load") the isotope into the scintillator. Should a $0\nu\beta\beta$ signal be observed, the scaling of the signal peak with isotope mass will be crucial in confirming the discovery.

In large liquid scintillator detectors, energy deposits from primary particles result in emission of light, which is then detected by photomultiplier tubes surrounding the scintillator volume. The number of detected photons is proportional to the energy deposited by the particle. Typically, there are $\mathcal{O}(10^3 - 10^4)$ photons emitted per 1 MeV of energy deposited in the scintillator [56, 57], which leads to high signal efficiency. Contrary to solid state detectors, these experiments have exceptional isotope mass scalability and modest energy resolution.

A currently operating $0\nu\beta\beta$ liquid scintillator detector is KamLAND-Zen [58]. It uses liquid scintillator loaded with ^{136}Xe -enriched xenon, enclosed in a nylon balloon. This volume containing the double β isotope is located in the centre of a larger balloon with 13 m diameter filled with xenon-free scintillator, which

acts as an active background veto. Approximately 2000 photomultiplier tubes, responsible for the detection of scintillator light, are mounted on a stainless steel spherical vessel outside both of the balloons in a non-scintillating oil. The KamLAND-Zen experiment achieved the current best limit on $T_{1/2}^{0\nu}$ of 1.07×10^{26} years [21]. This result was accomplished with 380 kg of enriched xenon. Since then, the collaboration has increased the inner balloon volume to double this amount and the data taking is ongoing. In the future, KamLAND2-Zen aims to have an improved energy resolution compared with its predecessor and will incorporate one tonne of enriched xenon [59].

The SNO+ experiment [60] will study $0\nu\beta\beta$ in ^{130}Te in a liquid scintillator detector, repurposed from the SNO experiment [61]. The search for this decay in SNO+ is the subject of the following chapters of this thesis.

2.4 Choice of Double Beta Isotope

As discussed in section 2.2, there are multiple naturally occurring even-even nuclei that can undergo double β decay. However, some are more suitable for a $0\nu\beta\beta$ search than others. The choice of the double β isotope is an important part of detector design and influences the sensitivity of the experiment to the $0\nu\beta\beta$ process. When selecting which isotope to use for this search, one has to consider the following factors:

- Firstly, high **Q-value** is desired. The phase-space factor from equation 2.12 scales with the decay endpoint and, hence, higher $Q_{\beta\beta}$ typically improves the chances of $0\nu\beta\beta$ detection by lowering the half-life of the process. Furthermore, many natural radioactive decays produce energies in the range of a few MeV, but the situation is especially bad below 2 MeV. With larger $Q_{\beta\beta}$, the ROI lies above many of these radioactive backgrounds. Finally, signal efficiency is generally better for larger energy releases.

- The **natural abundance** of an isotope gives the relative amount of that isotope in the naturally occurring element. It is advantageous to select a double β isotope with high natural abundance in order to maximise $N_{\beta\beta}$ for a given element mass. The abundances of the popular $0\nu\beta\beta$ isotopes are shown together with $Q_{\beta\beta}$ in Figure 2.5.
- Alternatively, elements can be enriched in the required isotope, but the **enrichment cost** tends to be high and considerably increases the overall **procurement cost**.
- Especially for experiments that plan to incorporate several hundreds of kilograms of a given element, **availability on world market** must also be taken into account. Ultimately, depending on the $\langle m_{\beta\beta} \rangle$ chosen by nature, it might be needed to design experiments with sensitivities reaching far into the normal hierarchy. This will likely require tonnes of an isotope and, therefore, the concern of availability is relevant for any experiment that might hope to scale up in the future.
- Long $2\nu\beta\beta$ **half-life** improves the experiment's sensitivity by reducing this background. This is especially important for detectors with worse energy resolutions. Figure 2.6 shows $T_{1/2}^{2\nu}$ for various candidate isotopes.
- The **nuclear factors**, $M_{0\nu}$ and $G_{0\nu}$, are specific for the candidate nuclei and determine the $0\nu\beta\beta$ half-life, with larger factors resulting in a better prospect of observation for the same $\langle m_{\beta\beta} \rangle$. Additionally, the **uncertainty on the nuclear matrix** is smaller for some isotopes compared to others.
- Any **technological limitations** must be considered. As discussed in section 2.3, some detector designs allow for variety of isotopes, whereas for others the choice is restricted.

As a result of the above criteria, the experimental efforts focus mostly on the following isotopes with $Q_{\beta\beta}$ larger than 2 MeV: ^{48}Ca , ^{76}Ge , ^{82}Se , ^{96}Zr , ^{100}Mo ,

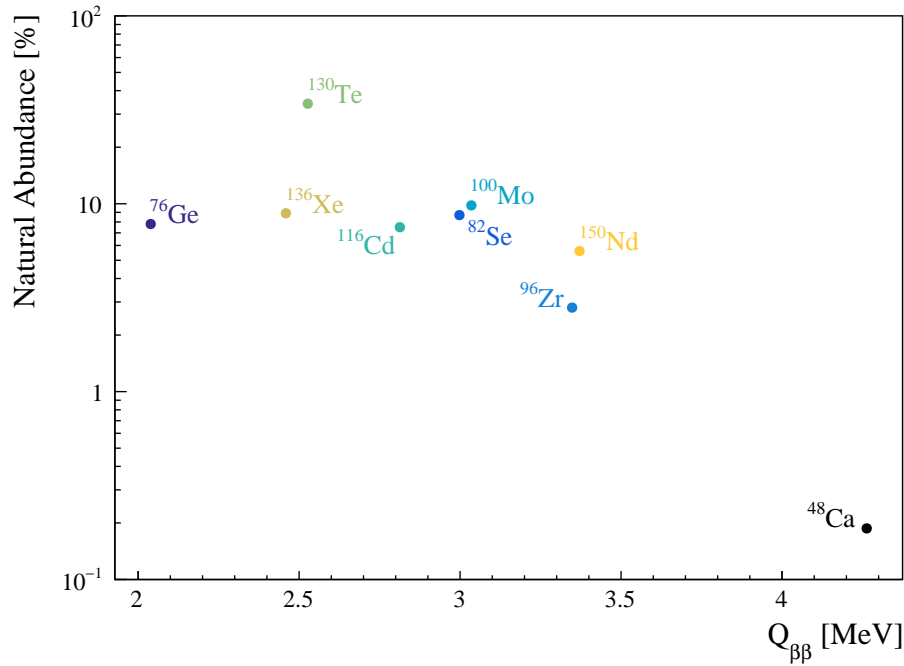


Figure 2.5: Double β decay Q -values as a function of natural abundance for selected double β nuclei. Data from [63] and [64].

^{116}Cd , ^{130}Te , ^{136}Xe and ^{150}Nd . As shown in Figure 2.5, ^{130}Te has by far the largest natural abundance (34%) out of the candidate nuclei, which mitigates the need of enrichment and makes it very cost effective. ^{130}Te also has the longest $T_{1/2}^{2\nu}$ after ^{136}Xe and ^{76}Ge , both of which are a lot less affordable, have lower $Q_{\beta\beta}$ and are not produced in high enough annual amounts to support a potential future multi-ton experiment [62].

For these reasons, ^{130}Te was selected as the double β isotope of choice for SNO+, the organic liquid scintillator experiment considered in this thesis. The method for loading natural tellurium into organic scintillator has been demonstrated and is discussed in Chapter 4. The SNO+ experiment is described in more detail in the following chapter.

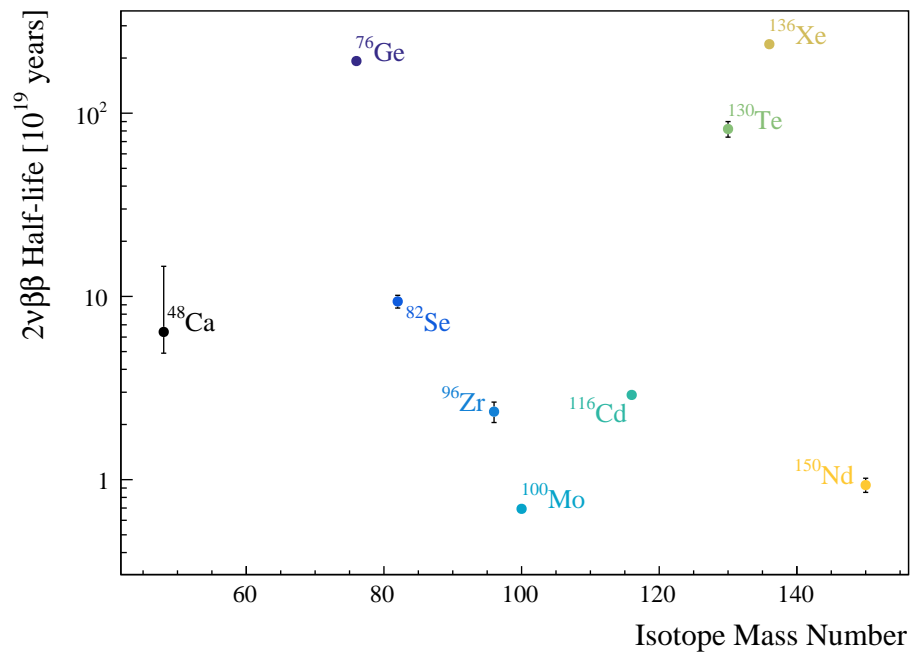


Figure 2.6: $2\nu\beta\beta$ half-lives for selected isotopes. Data from [7].

Chapter 3

The SNO+ Experiment

SNO+ is a kilo-tonne scale low background neutrino detector, located 2 km underground in an active nickel mine in Sudbury, Canada [60]. The experiment has 3 phases, distinguished by different material filling the main detector volume. Although its physics programme is diverse, the main goal of SNO+ is the search for $0\nu\beta\beta$. For this purpose, approximately 4 tonnes of natural tellurium will be deployed into a liquid scintillator volume using a loading procedure described in Chapter 4. SNO+ benefits from its size, allowing for large isotope exposure per year. The substantial size combined with low background rates gives SNO+ good expected sensitivity for $0\nu\beta\beta$. This chapter provides an overview of the detector hardware, software and calibration as well as the explanation of the various physics phases of the experiment.

3.1 Detector Overview

A schematic of the SNO+ detector is shown in Figure 3.1. Located in the centre, is a transparent spherical acrylic vessel (AV) with a radius of 6 m and a thickness of 5.5 cm. The physics analyses in SNO+ concern particle interactions occurring within this volume. The AV and its content are viewed by $\sim 10^4$ photomultiplier

tubes (PMTs) mounted on a stainless steel support structure (PSUP) with an average radius of 8.9 m. The outside of the AV, including the PMTs, is immersed in ultra-pure water (UPW) to shield the central detector region from cavity rock and PMT radioactivity.

Two sets of high purity Tensylon ropes are installed to hold the AV in place. The hold-up ropes (HUR), 19 mm in diameter, suspend from the top of the cavity and are hooked to ridges in the AV around its equator. The hold-down ropes (HDR) are 38 mm in diameter, wrap around the top of the vessel and are anchored to the cavity floor. The latter set has been installed in order to compensate for the buoyancy of the liquid scintillator relative to the surrounding water.

The barrel-shaped cavern and large hardware structures, such as the acrylic vessel, PSUP as well as the PMTs, were inherited from the SNO experiment [61], which was decommissioned in 2007.




Figure 3.1: Schematic of the SNO+ detector. Modified from [65].

The experiment is hosted in SNOLAB, a clean room CLASS 2000 laboratory, accessible from shaft 9 of VALE's Creighton mine. The deep underground loca-

tion provides SNO+ with rock overburden of 5890 ± 94 m.w.e. [66], minimising background from cosmic rays and their spallation products.

3.2 PMTs and Electronics

Interactions inside the SNO+ detector yield photons, which, depending on the phase, are either produced from Čerenkov or scintillation radiation. All information about events inside the detector come from detection of this light by the PMT array. For each triggered event, the PMT electronics output charge and timing data from each PMT that registered at least one photon.

3.2.1 PMTs

The bulk of SNO+ PMTs are 8" Hamamatsu R1408 SNO photomultiplier tubes [61]. The PSUP has space for 9522 PMTs. Over 75% of the ~ 800 PMTs that failed during the SNO data taking period were repaired and re-installed for SNO+, however, some were deemed unrepairable [65]. Furthermore, some of the PSUP PMT positions are not filled to provide space for the hold-down ropes. Finally, a small fraction of PMTs is continuously lost to high-voltage problems and at the rate of ~ 6.8 per month. At the time of writing, there were approximately 9200 fully functional PMTs facing towards the centre of the AV.

There are also 91 PMTs mounted on the outside of the PSUP pointed away from the AV. These outward looking PMTs serve as a veto for muons and other events originating in the external water.

A schematic of a SNO PMT is shown in Figure 3.2. The photoelectrons are created on the photocathode, which lies just underneath the round PMT glass envelope. These are then accelerated in the electric field permeating the vacuum inside the tubes and, subsequently, multiplied by a cascade of secondary emissions in a stack of 9 parallel dynodes. Finally, the photoelectrons generate a

measurable signal when they are collected on an anode at the base of the PMT.




Figure 3.2: A schematic of the Hamamatsu R1408 PMT as used in SNO+ [61]. The round PMT glass is shown on the top, the solid horizontal lines depict the dynode stack and the anode is the small dip near the bottom. The waterproof enclosure is also shown around the anode. Dimensions are in cm.

3.2.1.1 Transit time

The time between a photon arriving on the photocathode and the peak of the signal pulse registering on the anode is called the transit time and is defined by the path of the photoelectron within the PMT. A typical probability distribution of the transit time is shown in Figure 3.3. The main "prompt" peak (at 0 ns) corresponds to an ordinary path described above and its width is determined by the electron optics within the PMT geometry.

The two peaks later in time are formed by the late and double pulsing effects, which arise when the photoelectron is backscattered off the dynode stack. For late pulsing, this process is elastic and the cascade of secondary electrons is only created when the electric field returns the photoelectron to the dynodes $\sim 10 - 40$ ns later. In the case of double pulsing, scattering is inelastic and, hence, cascading occurs both times the photoelectron arrives on the stack. Therefore, these photoelectrons produce two pulses: one arriving at the same time as the

prompt pulse and one contributing to the broadening of transit time spread towards later times.

The earliest peak found at ~ -18 ns is from pre-pulsing. In this case, the incident photon is transmitted through the glass and photocathode and a photoelectron is created on the first dynode.




Figure 3.3: Transit time probability distribution for a R1408 PMT plotted using values saved in the SNO+ simulation software. Originally measured by the SNO Collaboration.

3.2.1.2 Wavelength efficiency

The total PMT signal efficiency describes the probability of an incident photon of a given wavelength to produce a signal on the anode. Quantum and collection efficiencies are the two contributions to the overall signal probability. The former gives the ratio of photoelectrons leaving the photocathode to incident photons. The latter is defined as the ratio of photoelectrons reaching the first dynode to the photoelectrons leaving the photocathode. Figure 3.4 shows the measurement of the combined effect for SNO PMTs.




Figure 3.4: Measured wavelength efficiencies for 4 R1408 PMTs [67]. Each symbol corresponds to a different PMT.

3.2.1.3 PMT Concentrators

Each inward looking PMT is also equipped with a truncated Winston cone-shaped reflective concentrator with an outer diameter of 27 cm [68]. These increase the effective photomultiplier coverage of the detector by focusing photons on the active area of the photocathode. At the start of the SNO experiment data taking, the concentrators ensured an effective coverage of 54% [69]. Since then, the reflectivity of the concentrators has degraded and the number is now closer to $\sim 50\%$. The PMTs and their concentrators are housed in plastic hexagonal boxes, which enable their installation in groups to the PSUP. Figure 3.5 shows the complete PMT unit including the concentrator and housing.

3.2.2 Electronics

The electronics are responsible for the read out, analogue processing and digitalisation of the PMT signals. A trigger system is responsible for distinguishing which signals from the detector are likely to correspond to real physics events and are worth saving. In order to provide enough information about the physics




Figure 3.5: Drawing of a PMT in its housing including a concentrator [61]. Dimensions are in cm.

inside detector for the purpose of data analysis, the electronics must produce nanosecond-level timing resolution and be able to handle sustained frequencies of several kHz.

Every PMT is coupled to a waterproof coaxial cable, used to both supply the high voltage (HV) and send the anode pulse to the front-end electronics located on an access deck above the SNO+ cavity. There, it connects to a PMT Interface Card (PMTIC), which controls the HV for the corresponding PMTs. One PMTIC has 32 channels, each potentially connected to a PMT. From the PMTIC, the anode signal continues to a Front End Card (FEC), where it is processed.

To distinguish between real signals and electronic noise, the pulse enters a discriminator on the FEC, which only fires if the pulse amplitude is higher than the discriminator threshold. The threshold is usually set to around 0.25 of an average photoelectron charge for the relevant channel [65]. This value is determined from a single photoelectron charge spectrum as the value where the signal is more likely to be caused by an incident photon rather than noise. As a result, about a quarter of genuine PMT signals are falsely identified as noise. A PMT is considered hit if the discriminator threshold has been crossed.

Sets of 16 FECs and PMTICs make up "crates", which are located in electronics racks on the deck above the detector. There are 19 crates altogether, inside 11 racks, creating capacity for 9728 channels. The discriminator pulses are first summed within each crate and these crate sums are then passed to seven MTC/A (Master Trigger Card/Analog) boards, each corresponding to a particular physics trigger type. Each MTC/A board performs an analog sum of the crate signals and issues a trigger if that sum exceeds some pre-determined threshold. As an example, one MTC/A is responsible for the NHIT100 trigger, which is one of the most used for physics analyses. The amplitude of the NHIT100 MTC/A pulse is proportional to the number of hit PMTs within 100 ns.

Finally, MTC/A trigger pulses are sent to the MTC/D (Master Trigger Card/Digital). MTC/D can be programmed to acknowledge or ignore (i.e. mask in or out) individual trigger types. Provided the received trigger pulse is masked in, the Global Trigger (GT) is sent back to the FECs and the detector is triggered. When this occurs, all PMT data are collected from a period 180 ns before the GT happened and 220 ns after. This 400 ns period is called the "event window", as it would typically contain information about one interaction event inside the detector. Conversely, the word "event" is henceforth used to refer to detector information recorded within a single event window.

If required, it is also possible to force the MTC/D to generate a global trigger externally, regardless of the PMT channels. The EXTA (External Asynchronous) trigger is issued by many of the calibration sources described in section 3.4 and is used to prompt the PMT data acquisition during calibration campaigns.

There are four numbers output for each PMT hit. The most important is the TAC (Time to Analog Converter), which gives the time between the hit (PMT crossing its threshold) and the GT reaching its PMTIC. Hence, the TAC indicates the relative time within the event window that a particular PMT fired, which is important for event reconstruction. The remaining three numbers associated

with a PMT hit all relate to charge, but differ in the gain and the time over which the charge was integrated; there is QHS (High gain Short integration), QHL (High gain Long integration) and QLX (Low gain Long integration).

3.3 SNO+ Target and Phases


The central target for physics interactions in the SNO+ experiment, is the liquid inside the AV. One of the advantages of SNO+ is that this liquid can be exchanged or modified relatively easily, which gives the potential to study a wide range of physics topics. Furthermore, the detector model and many experimental backgrounds for the analyses, such as the search for $0\nu\beta\beta$, can be studied prior to introducing the isotope.

Each time the inner AV material is modified, start of a new phase of the experiment is defined. This section gives an overview of the planned phases and describes each of their purposes.



3.3.1 Water Phase

The HV in the detector was first ramped up in December 2016 and the water phase officially started in May 2017 after an initial commissioning period. In this phase, the AV was filled with 905 t of ultra-pure water.

The main physics analysis carried out in this phase was the search for invisible nucleon decay. In this hypothetical process [70–72], a nucleon decays into neutrinos or other particles that leave the detector without any direct energy deposition (*inv*). However, if the decay occurs in an oxygen atom in the SNO+ water, the nucleus is left in an excited state and subsequently decays into its ground state. For neutron decay:



and proton decay:



In more than 45 % of cases, the de-excitation γ s have energies in excess of 6 MeV, which are easily detectable in SNO+ [73]. During almost 8 months of data taking, lower limits on the neutron and proton lifetimes were set to be 2.5×10^{29} and 3.6×10^{29} years, respectively [74].

Furthermore, the flux of ^8B solar neutrinos was measured [75]. This result was found to be consistent with the measurements by other experiments [10,76] and demonstrated low levels of backgrounds. The background measurements are further discussed in Chapter 5.

The optical properties of the external water and the AV, the response of the PMTs and the performance of the data acquisition system were also measured over this period. Finally, the backgrounds originating from outside of the AV, such as the support ropes, external water and PMTs, were measured [74].

3.3.2 Unloaded Scintillator Phase

One main factor that can impact an experiment's sensitivity to $0\nu\beta\beta$ is the level of radioactive background contamination from U/Th chains in the detector volume. Organic liquid scintillators (LS) can have very low amounts of backgrounds as most impurities are ionic and, hence, do not dissolve in non-polar organic liquids. This has been ultimately demonstrated by the Borexino experiment, which reached contamination levels of less than 10^{-18} g/g for both ^{238}U and ^{232}Th chains [55]. Furthermore, LS is easily and economically scalable. For these reasons, SNO+ will use a LS as a solvent and detection target for the $0\nu\beta\beta$ search.

The radiopurity of the SNO+ scintillator will be first measured in the unloaded scintillator phase with no double β isotope present. The LS is based on linear alkylbenzene (LAB), a long chain aromatic molecule, with

2,5-diphenyloxazole (PPO) as the primary fluor [60]. PPO reduces self-absorption by LAB and also shifts the emitted photon energies to a spectral range where the SNO+ PMTs have higher sensitivity [77]. Properties of the scintillator mixture are discussed in detail in Chapter 4.

In order to replace the ultra-pure water inside the AV with scintillator, the LAB+PPO is filled from the top via the neck, with water being removed by a pipe at the bottom of the vessel. The scintillator mixture has lower density than the ultra-pure water, which allows it to float on top. Moreover, the two liquids are immiscible, which results in clear interface between them during filling as shown in Figure 3.6. This displacement approach also minimises air (and, hence, radon) contact with the AV.




Figure 3.6: Photo of the neck and top part of the AV during scintillator fill. Taken on 24th July 2019 using one of the under-water cameras mounted on the PSUP. The mass of the scintillator in this photo is ~ 20 tonnes. Photo modified from [78].

Filling of LAB began in October 2018. Once finished, the unloaded scintillator phase will commence and last for a few months. During this time, the optical model and detector response will be verified, along with the evaluation of the radioactive background levels of the scintillator mixture.

The physics goals in this phase include the observation of low energy solar

neutrinos, reactor and geo anti-neutrinos and potential supernova neutrinos.

3.3.3 Tellurium Phase

The $0\nu\beta\beta$ search will be conducted in the tellurium phase. Four tonnes of natural tellurium will be gradually loaded into the LS for this phase. The scintillator from the AV will be recirculated through a processing plant underground, where it will be mixed with the tellurium compound. The loading technique is detailed in Chapter 4.

3.4 Detector Calibration

A good understanding of the detector properties is needed in order to translate the PMT information into physical variables, such as the time and energy of the interaction that caused the detector to trigger. The purpose of calibration is to measure these properties and, hence, constrain all the free parameters in the detector model.

The main quantities to calibrate are summarised below:

- The time of each PMT hit has to be corrected for **channel-dependent timing offset** due to variable lengths of the PMT cables. Additionally, smaller amplitude pulses take longer to cross discriminator thresholds than larger pulses. This is called the **time walk effect**, which must also be measured and taken into account during the hit time calculation.
- The PMT signal efficiency includes the quantum and collection efficiencies, as described in section 3.2.1. To determine the probability of detecting a single photon of a given wavelength, the signal efficiency has to be multiplied by the individual **channel efficiencies** due to the discriminator threshold settings.

- The **angular response** of the PMTs and their concentrators modifies the detection probability for photons incident from different angles.
- The **scattering and absorption lengths** of the liquid contained in the AV, external water and the AV itself are needed to model the propagation of light through the detector.
- The reflectivities of the various materials in the detector also contribute to the optical model. Most importantly, the **concentrator reflectivity** determines the effective photocathode coverage of the detector. Furthermore, the differences in refractive indices between scintillator, acrylic and water causes reflection/refraction at boundaries. These need to be taken into account when calculating the time-of-flight of photons from their origin to the PMT.
- The scintillator **light yield** determines how many photons are emitted per 1 MeV of energy deposited in the scintillator. The light yield therefore plays a crucial role in energy reconstruction. Monitoring of the light yield over time is needed to verify its stability.
- **Time response of the scintillator** is important for position reconstruction as well as particle identification as it varies with particle type.

Some of these detector parameters, such as the scintillator properties or PMT response, can be determined ex-situ to provide an initial measurement. However, a number of the properties can change with time. For example, the PMT concentrator reflectivity was measured on a benchtop before installation, but has changed since then due to ageing and their immersion in UPW. Similarly, it is impossible to precisely mimic the detector environment for a small sample of the scintillator over multiple years. Therefore, in-situ calibration is needed, which is achieved using optical and radioactive sources inside the detector.

3.4.1 Calibration Using Radioactive Sources

Deployed radioactive sources can be used to introduce particles of known type and energy into the detector in order to calibrate the detector response to primary particle interactions. In particular, the energy scale and resolution, as well as systematic errors on reconstructed quantities can be obtained using this method. For this purpose, sources are lowered into the detector through the neck using a system of ropes, with each source enclosed in a leak-tight housing to prevent contamination.

In the water phase, the systematic uncertainties for the nucleon decay and solar flux analyses were obtained using the ^{16}N source. In this case, the β decay of ^{16}N produces a γ at 6.1 MeV, which escapes the source container and interacts in the detector target material. The β from the decay interacts inside the container, where it is detected in a thin layer of plastic scintillator coupled to a calibration PMT. This signal is used to tag the γ events. For future phases, sources of other energies will be deployed to reduce the systematic errors across the energy spectrum. A list of radioactive sources considered for future deployment in scintillator is shown in Table 3.1.

3.4.2 Optical Calibration

In optical calibration, LEDs and lasers are coupled to optical fibres to direct photons of fixed wavelengths into the detector. Optical sources are the ideal tool for measuring the PMT response and optical properties of the detector target. Furthermore, these systems can trigger the detector asynchronously, making it possible to measure timing offsets. The two optical calibration systems for SNO+ are described below.

Source	Particle type	Visible energy [MeV]	Calibrated quantity
AmBe	n, γ	2.2, 4.4	neutron capture response
^8Li	optical photons	-	Čerenkov light response
^{57}Co	γ	0.1	energy reconstruction
^{137}Cs	γ	0.7	energy reconstruction
^{46}Sc	γ (2x)	0.9, 1.1	energy reconstruction
^{48}Sc	γ (3x)	1.0, 1.0, 1.3	energy reconstruction
^{16}N	γ	6.1	energy reconstruction

Table 3.1: Radioactive calibration sources in SNO+, available or in preparation. The second column indicates the type of particles that deposit energy in the detector material, i.e. particles that do not escape the source container are not included.

Laserball

The laserball is a quartz light-diffusing sphere of diameter 10.9 cm and was originally designed for the SNO experiment [79]. The sphere is filled with silicone gel containing small hollow glass beads. Light from a short-pulse nitrogen dye laser is brought into the sphere via an optical fibre, where it is reflected and refracted by the beads. This results in an approximately isotropic distribution of light coming from the laserball. The intrinsic wavelength (337 nm) can be shifted to five higher values (369, 385, 420, 505 and 619 nm) using laser dyes.

The laserball can be deployed either into the liquid inside the AV or into the water region between the AV and the PSUP. It is deployed in the same way as the radioactive sources, via a set of ropes. The laserball is one of the primary calibration sources used for the optical model of the detector. By fitting the laserball data taken over a range of wavelengths and at various positions in the detector, it is possible to extract the attenuation parameters for the different detector regions [80]. Furthermore, the laserball is used to determine the individual PMT/channel efficiencies and, using asynchronous triggering of the detector, the

channel timing offsets and time walk.

ELLIE System

The radiopurity requirements of SNO+ are more stringent when compared to SNO. Every effort is made to mitigate the risk of contaminating the detector material by the introduction of deployed instrumentation. Even though extensive cleanliness protocols for the calibration sources are in place, the easiest way to avoid contamination is to reduce the frequency with which the sources are immersed into the detector. For this purpose, the Embedded LED/Laser Light Injection Entity (ELLIE) was designed for SNO+ [81]. This calibration system consists of optical fibres permanently mounted at various points of the PSUP, which inject LED or laser light across the detector volume at different angles.

The ELLIE system has three modules:

- The timing module (TELLIE) injects wide angled LED beams of a single wavelength from 91 different positions on the PSUP to provide timing and gain calibration.
- The scattering module (SMELLIE) consists of 15 fibres delivering collimated beams of laser light at 5 wavelengths to measure scattering properties of the materials inside the detector.
- The attenuation module (AMELLIE) shines wide angled beams of LED light at multiple wavelengths to measure attenuation lengths.

With this set up, ELLIE calibration campaigns do not require any instrumentation to be moved into the detector and the downtime is minimised. In addition, this easy operation allows ELLIE to take data frequently and, hence, monitor the stability of the measured variables over time.

3.5 Software

Reactor Analysis Tool (RAT) is used by SNO+ for both data analysis and simulation of physics events in the detector. It is a C++ framework that makes use of Geant4 [82] and GLG4sim [83] libraries for detailed Monte Carlo simulation (MC). Geant4 is used to model the detector geometry including all the components. Various particle generators enable the simulation of $0\nu\beta\beta$, solar and geo neutrinos, anti-neutrinos, muons and the decays of the expected radioactive background isotopes. Optical photons are generated via Čerenkov and scintillation processes from GLG4sim. The propagation of photons through the detector is fully modelled including reflection, scattering, absorption and re-emission. The effects of PMT response and front-end electronics are also simulated. Furthermore, analysis algorithms, such as event reconstruction, are incorporated into the framework. Finally, an interface with ROOT is provided for data storage and handling.

Chapter 4

Tellurium-Loaded Scintillator

Tellurium-130 was selected to serve as the double β isotope for the $0\nu\beta\beta$ search in SNO+. This chapter focuses on the method to combine tellurium with the liquid scintillator in the AV, a process generally referred to as tellurium-loading. In particular, increasing the stability and light yield of the resulting tellurium-loaded scintillator through the addition of a stabilising compound, N,N-dimethyldodecylamine (DDA), is investigated here.

The topics concerning the scintillation process, as well as the composition of the scintillator prior to tellurium loading, are also discussed at the beginning of this chapter.

4.1 Scintillation Process

During the scintillation process, light is emitted as a result of ionising radiation passing through a material. The molecules of the scintillator are excited by the through-going particles and the subsequent relaxation to the ground state is responsible for the photon emission. Therefore, scintillation light has a characteristic wavelength profile that is dictated by the chemical structure of the given scintillator.




Figure 4.1: Schematic representation of π electrons in benzene. Carbon atoms are represented as red spheres. σ bonds are depicted as black lines and hydrogen atoms are omitted for simplicity. On the left, six individual p orbitals are shown in blue. These π orbitals overlap and form clouds of delocalised charge as shown on the right.

Scintillation occurs in many types of materials, including organic and inorganic, solid, liquid and gaseous, crystalline and amorphous [84]. For the purpose of large-scale neutrino detectors, organic liquid scintillators are a convenient choice due to their ease of purification, scalability and their low cost. These are typically hydrocarbon molecules that contain one or more aromatic benzene rings. Benzene has planar geometry and molecular orbitals with sp^2 hybridisation, which connect each carbon to two adjacent carbons and one hydrogen at 120° . The electrons involved in bonding the atoms through the overlap of these molecular orbitals in the plane of the ring are the σ electrons. Three electrons per carbon contribute to the σ bonds. The last electron of each carbon is left in a p orbital oriented perpendicularly to the plane. The p orbitals of neighbouring carbons overlap and form secondary π bonds. As a result, there is a cloud of delocalised charge above and below the ring, meaning that the electrons are shared between all the carbon atoms, as shown in Figure 4.1. Furthermore, molecules with delocalised charge can have energy levels with separation corresponding to energies of optical photons. It is the excitation of these π orbitals that is responsible for scintillation.

In particle detectors, the excitation of the π cloud can happen through elastic scattering or ionisation followed by recombination into the excited state [85]. The subsequent de-excitation results in the emission of isotropic scintillation light. Only spin singlet states can be created via direct excitation from elastic scattering because the ground state is also a spin singlet. With ionisation and recombination, however, both singlet and triplet states are allowed. While the former can de-excite immediately ($\mathcal{O}(10^{-8}$ s)), the triplet states are metastable, which results in slower scintillation light emission. It follows that a higher ratio of ionisation to direct excitation (and, hence, heavier ionising particles) leads to a larger slow component of the overall emission timing profile.

Finally, in order to prevent the scintillation photons from being absorbed by the material again, there must be a shift between the absorbed and emitted energy spectra. To some extent, this is taken care of by vibrational relaxation: the molecule can be excited to any vibrational state of the electronic state, which is then (non-radiatively) relaxed to the lowest one and only decays radiatively from there. Therefore, the energy of the photon emitted from this state is lower than the absorbed energy and, hence, cannot cause the same excitation.

4.1.1 Fluors

Fluors (also called fluorophores) are compounds capable of absorbing energy and re-emitting it as a photon with wavelength corresponding to a lower energy. Therefore, the addition of fluors can also be used to prevent self-absorption of scintillation light [84]. Moreover, de-excitation of scintillator by a non-radiative process is undesirable, but also possible. The fraction of excitations resulting in radiative emission is called the quantum yield of the scintillator. Fluors can have higher quantum yield than the solvent and, hence, their addition can improve the light output of the overall mixture.

Primary fluors are those to which energy can be transferred directly from the

solvent scintillator. The typical concentrations of primary fluors is g/L, where the energy transfer from scintillator to fluor is mostly non-radiative. Secondary fluors, to which energy is usually transferred radiatively from the primary fluor, can also be used to further adjust the resulting emission spectrum to the detector requirements.

4.1.2 Light Yield

One important property of a scintillator mixture is its light yield (LY), defined as the number of emitted photons per quanta of energy absorbed by the scintillator. It is described by the Birks' law:

$$\frac{dL}{dx} = L_0 \frac{dE/dx}{1 + k_B dE/dx}, \quad (4.1)$$

where L is the light yield, L_0 is its value at a specific low ionisation density and dE/dx is the energy deposit by the particle per unit length [84]. The second term in the denominator describes the fact that for heavier ionisation particles the light output for the given energy is decreased due to "ionisation quenching" (this is distinct from chemical quenching discussed in section 4.4.2). Birks ascribes this to the interactions between excited molecules leading to non-radiative de-excitations. The effect is empirically described by a larger value of the Birks' constant, k_B , for heavier particles compared to light radiation for a given scintillator.

In SNO+, PMT hits are detected rather than individual photons. Therefore, LY is often given in hits/MeV, despite the fact that the relationship between photons and hits is not perfectly linear (see section 5.2.2).

4.2 Scintillator in SNO+

The main component of the scintillator mixture in SNO+ is linear alkylbenzene (LAB), which is used as the primary solvent. LAB is liquid at room temperature and has a long hydrocarbon chain attached to an aromatic benzene ring, as shown in Figure 4.2. 2,5-Diphenyloxazole (PPO) will be used as the primary fluor. The concentration of PPO has been chosen to be 2 g/L to provide a good scintillator LY (~ 600 PMT hits/MeV) for a reasonable cost.




Figure 4.2: The chemical structures of LAB, PPO and bisMSB.

Finally, 1,4-bis(2-methylstyryl)benzene (bisMSB) will be used as the secondary fluor in the tellurium phase. The exact concentration is yet to be optimised, but even several mg/L of bisMSB shifts the wavelength of the scintillation light to a region closer to the PMT efficiency maximum. It can also help to move the emission spectrum away from potential absorption associated with the tellurium loading method. The energy transfer to bisMSB is radiative, and therefore, the intrinsic light yield, i.e. the number of initially produced photons, is not affected.

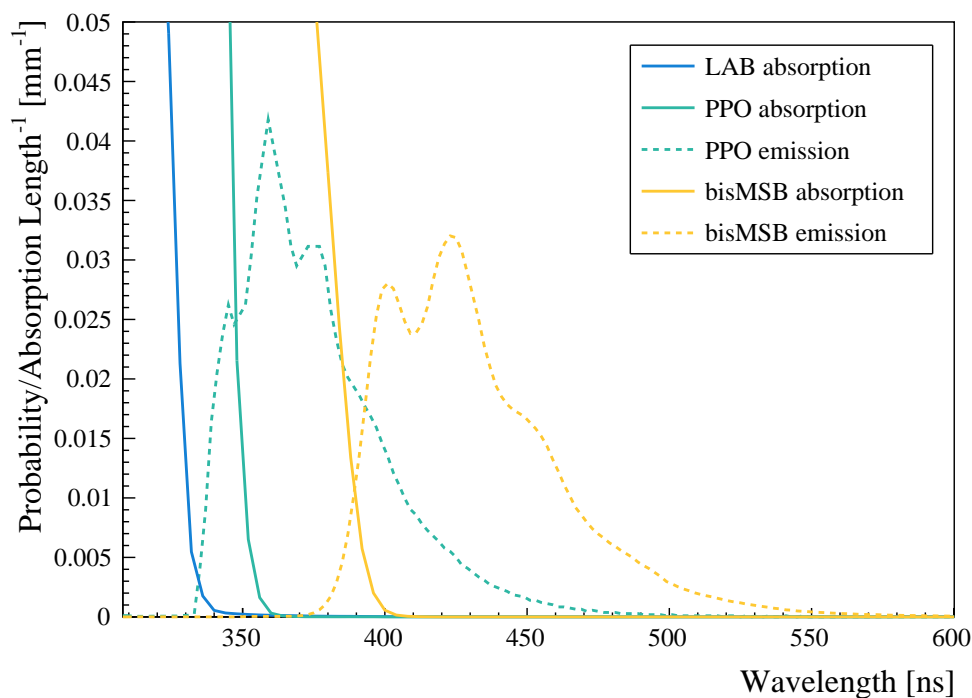


Figure 4.3: The inverse absorption lengths of LAB (blue full line), PPO (green full line), and bisMSB (orange full line) as a function of wavelength as used in RAT. Originally measured in [86,87]. The emission spectrum of the LAB+PPO is modelled to be that of PPO alone (green dashed line). The emission spectrum of bisMSB is shown in orange dashed line.

The number of photons reaching the PMTs is reduced according to the quantum yield of bisMSB (~ 0.95), but the overall detection efficiency is improved. As the addition of bisMSB is not expected to change the intrinsic light yield or the scintillator stability, no bisMSB was used for the studies shown in this chapter for simpler sample preparation. However, the effects of this secondary fluor have been tested and its addition is assumed in the work shown in subsequent chapters.

The absorption coefficients of the above components are shown in Figure 4.3. As the light absorbed by LAB is non-radiatively transferred to the fluor, the emission spectrum of the LAB+PPO mixture is identical to that of PPO, which is also plotted.

4.3 Requirements on Tellurium Loaded Scintillator

In order to search for $0\nu\beta\beta$ in tellurium-130, natural tellurium will be introduced into the LAB+PPO mixture. The resulting loaded scintillator cocktail has to comply with the following requirements:

- In order to maintain good energy resolution, the loaded scintillator must retain **high light yield** of $\mathcal{O}(10^2 \text{ PMT hits/MeV})$. In general, the higher the light yield, the better the experimental sensitivity is to $0\nu\beta\beta$.
- The tellurium phase considered in this work will last for multiple years (at least three and up to five). Therefore, **long-term stability** of the mixture of at least five years and **compatibility** with materials inside the detector is required.
- The loading should not hinder **good optical properties** of the scintillator. For example, it should not significantly increase scattering nor introduce absorption lines into the wavelength region where the PMTs are most efficient.
- Any material put inside the SNO+ detector has to have an **excellent radiopurity** to allow for rare event searches. The target contamination level of the Te-loaded scintillator is $\mathcal{O}(10^{-15} \text{ g/g})$ for both ^{238}U and ^{232}Th chain contaminants [88].
- Due to its location in an active mine, all chemicals used in SNO+ have to satisfy stringent **safety requirements**.
- **Affordable cost** of the all chemical compounds used for the preparation of the final scintillator cocktail is required.

4.4 Tellurium Loading

Inorganic tellurium compounds are generally immiscible in organic solvents such as LAB and, unfortunately, typical organic compounds tend not to fulfil the above criteria on several counts. Hence a new compound, tellurium butanediol (TeBD), had to be synthesised for SNO+ [89]. TeBD can be simply mixed with LAB in any required concentration, though light yield limits the practical loading to a few percent level. For the tellurium phase of SNO+, the target is 0.5% of natural tellurium by weight.

4.4.1 Tellurium Butanediol Synthesis

The synthesis of TeBD involves two main reactants: telluric acid (TeA), an inorganic molecule with a Te atom surrounded by six hydroxyl groups, and 1,2-butanediol (BD), a small organic compound with two hydroxyl groups and a short tail. Telluric acid is a white powder at room temperature and is first dissolved in water. 1,2-Butanediol is a colourless liquid and is poured directly into the aqueous solution of TeA.

Nuclear magnetic resonance (NMR) and mass spectrometry measurements of the solution after mixing the reactants show that the BD attaches to the TeA at the hydroxyl group positions as shown in Figure 4.4 [90]. There are six hydroxyl groups on the TeA and, hence, space for three BD molecules, each attaching with two bonds. However, a compound with such a configuration has not been observed. Instead, one or two (as in the figure) BD molecules are attached per Te nucleus at this stage. These molecules that contain one Te atom and at least one BD substituted OH group are the precursors of TeBD.

The precursors are still not soluble in organic solvent, likely due to the number of exposed OH groups. Furthermore, the synthesis of these compounds happens in aqueous environment and more water is created in the process. Not only is the water itself immiscible with LAB, its presence also makes it possible




Figure 4.4: Condensation reaction of telluric acid with 1,2-butanediol giving a TeBD precursor molecule with two substituted BDs.

for the reaction in Figure 4.4 to proceed in the backward direction. Therefore, if the reactants are mixed at room temperature, the condensation reaction is in equilibrium with hydrolysis.

To remove the water, drive the reaction towards condensation, and obtain a LAB soluble compound, the mixture is sparged with nitrogen gas, heated to 70 – 80°C and maintained at this temperature until majority of water is evaporated. The heat also promotes further BD attachments and polymerisation of the precursors into molecules with more than one Te atom. Figure 4.5 shows some of the possible chemical structures of the resulting molecules. NMR measurements show that these multiple forms are present simultaneously. TeBD complex is a term describing the mixture of these molecules, dominated by forms with at least two Te atoms surrounded by short organic hydrocarbon tails. It contains fewer OH groups per Te atom compared to the precursors. As a result, the TeBD complex is soluble in LAB.

4.4.2 TeBD Loading Effect on Scintillator LY

The TeBD complex is a transparent, highly viscous liquid. It does not have any absorption lines in the visible wavelength range and its addition does not significantly worsen the scattering properties of the scintillator [86].

However, the tellurium loading quenches the scintillator light to some extent, which results in an overall decrease in light yield compared to unloaded scin-




Figure 4.5: Examples of possible structures of molecules inside the TeBD complex. Forms with more than two Te atoms might also be present at lower concentration.

tillator. The term "fluorescence quenching" is commonly used to describe any process that results in the decrease of LY of scintillators. Quenching can arise due to molecular and electronic re-arrangements, transfer of energy between the molecules, collisions of the molecules or preferential excitation to triplet states in the presence of a heavy atom. In this case, it can be for example a result of the large electronegativity of the TeBD complex, which partially immobilises the π electrons in the aromatic rings, preventing the creation of scintillation light. The overall decrease in scintillation light is likely a result of a combination of the possible quenching processes and is a subject of ongoing studies.

This LY decrease due to Te loading is demonstrated in Figure 4.6, which shows a comparison of unloaded LAB + PPO and 0.5 % Te samples measured by a LY monitoring setup in Oxford. The setup uses a ^{90}Sr source to produce β particles, which pass through a cylindrical cuvette with outer diameter of 1 cm containing 3 g of the scintillator sample. The scintillator light is detected by

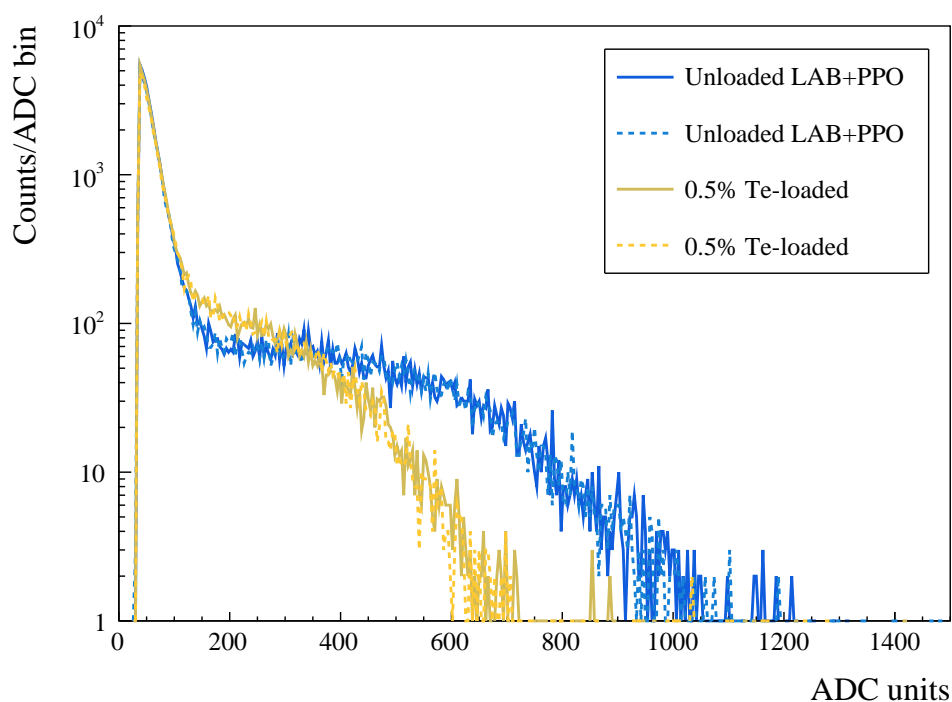


Figure 4.6: The charge intensity spectra of LAB + 2 g/L PPO with no (blue) and 0.5 % (orange) tellurium added in the form of TeBD. The analog-to-digital converter (ADC) units are arbitrary, but sufficient for relative comparison between samples. Two different samples (dashed and solid line) of the same concentration are plotted to show sample to sample variation.

a Hamamatsu H11432 PMT and a spectrum of collected charge is output by a multichannel analyser in arbitrary analog-to-digital converter (ADC) units. Data is recorded for 5 min for each sample. A photo and a block diagram of the LY measuring setup are shown in Figures 4.7 and 4.8, respectively.

The energy deposited in the scintillator sample depends on the β particle's energy and length of trajectory through the material. Therefore, both the underlying energy deposit and the recorded charge spectrum are continuous distributions. The shape of the charge spectrum depends on the full geometry of the setup as well as homogeneity of the sample itself. However, regardless of the shape, the endpoint of the spectrum corresponds to the maximum deposited energy during the data taking window. Thus, with identical exposure time for each sample, the endpoint is a reliable measure for relative light yield compar-




Figure 4.7: A photo of the LY monitoring set-up. The ^{90}Sr source is located below the cuvette stage and shielded from all sides apart from ~ 4 mm aperture on which the vial containing the scintillator mixture is centred. A monitoring PMT is fixed ~ 2 cm from the cuvette.

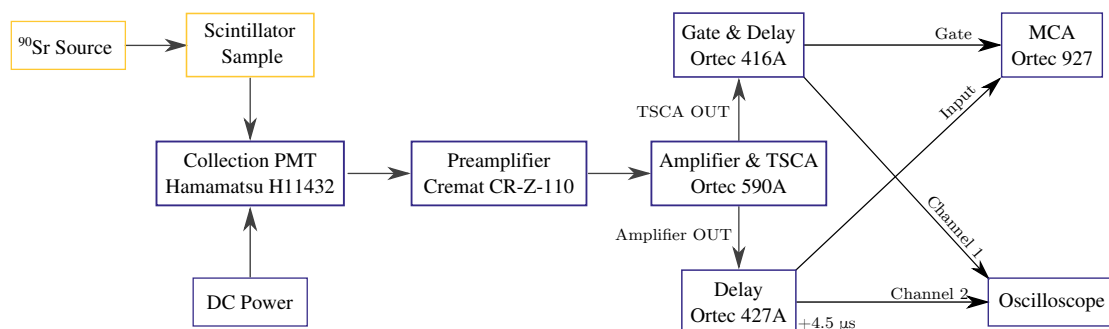


Figure 4.8: A block diagram of the LY monitoring set-up [91]. The signal from the PMT is first preamplified and then sent to an amplifier with a Timing Single Channel Analyzer (TSCA) unit, which produces a shaped amplified pulse (Amplifier OUT) together with a square pulse (TSCA OUT). The square pulse is made longer and used as a gate to initialise a 8k channel Multichannel Analyser (MCA). The shaped pulse is delayed in order to arrive on the MCA after the gate, where a spectrum of collected charge in ADC units is read out using the MAESTRO software.

ison between the samples.

A non-parametric estimation of the endpoint is employed throughout this work, taking the last 40 counts in the charge spectra. This method is robust against spectrum shape changes due to sample heterogeneity and has been tested to yield consistent results with absolute LY measurements at University of Pennsylvania and elsewhere [92].

The spectra shown in Figure 4.6 reveal a $\sim 40\%$ decrease in light yield upon loading the LAB+PPO mixture with 0.5% Te. The comparative LY measurements are also converted to PMT hits inside the SNO+ detector for various loading concentrations in Figure 4.9. The PMT hits for the unloaded scintillator were obtained from a RAT 6.15.6 simulation of 1 MeV electrons in the centre of the detector. The simulation used unloaded LY measured in fully calibrated setup described in [86]. The number of hit PMTs from 500 simulated electrons was confirmed to follow a Poisson distribution and therefore, Poisson error bars are shown in the figure. To plot the LY as a function of Te-loading, the PMT hits for the unloaded scintillator from simulation were then scaled down based on the relative LY measurement in Oxford for the given tellurium concentration. It is shown that the light yield falls significantly as more TeBD is added.

As the energy resolution in SNO+ improves with higher number of PMT hits, the sensitivity to $0\nu\beta\beta$ half-life increases with the exposure to the isotope (in kg-year) while also increasing with the LY. Therefore, the quenching limits the loading and affects the potential of the experiment.

4.4.3 TeBD Loading Effect on Scintillator Stability

In addition to impacting the light output, tellurium loading can also affect the scintillator stability. While scintillator mixtures with higher ($>10\%$) Te concentrations appear to be stable over long periods of time, deterioration of some samples with lower loading was observed in the form of flake formation. This




Figure 4.9: Effect of loading tellurium in the form of TeBD into LAB+2 g/L PPO on light yield. PMT hits/MeV were obtained from a SNO+ RAT simulation with unloaded scintillator and scaled according to LY measurements in Oxford. Two samples synthesised, mixed, and measured on different days are shown for each loading (blue squares and green circles) for sample variation estimation. Poisson error bars reflect statistical fluctuations.

"crash" is shown in Figure 4.10.

It was concluded by the SNO+ collaboration that the crashing is a result of Te-loaded scintillator being susceptible to water exposure. This is of a special concern in SNO+, because the scintillator will fill the inner volume of the SNO+ AV directly succeeding currently present water.

4.5 Amine Addition

It was proposed by a group at the Brookhaven National Laboratory (BNL) that a long chain amine could reduce quenching and potentially improve stability. Positive results were shown with octylamine at BNL and elsewhere [93]. However, octylamine is classified as severely toxic and flammable with a flash point of





Figure 4.10: Photos of uncompromised 0.5 % Te-loaded scintillator sample on the left and an example of flake formation on the right.

only 60°C [94]. Therefore, the deployment of octylamine underground and thus in SNO+ is unfeasible due to safety restrictions and an alternative was sought.

A variety of amines has been tested and the list of them is provided in Appendix B. In the end, *N,N*-Dimethyldodecylamine (DDA) was selected due to its non-biogenic origin, good miscibility with LAB, LY boost, and stability improvement. Moreover, it has a flash point of 118°C and lower toxicity than octylamine [95].

DDA is a tertiary amine with a long hydrocarbon chain and its structure is shown in Figure 4.11. The effects of the DDA addition are discussed below.



N,N-Dimethyldodecylamine

Figure 4.11: The chemical structure of *N,N*-dimethyldodecylamine (DDA).




Figure 4.12: Effect of adding DDA to 1% Te-loaded LAB + 2 g/L PPO on light yield. PMT hits/MeV were obtained from a SNO+ RAT simulation with unloaded scintillator and adjusted according to comparative LY measurements in Oxford. Duplicate samples (blue squares and green circles) are shown for each concentration for sample variation estimation. Poisson error bars are shown.

4.5.1 LY Improvement with DDA

Figure 4.12 shows the LY comparison of 1% Te-loaded scintillator sample with and without the addition of DDA at various DDA:Te molar ratios (MR). The measurements were obtained using the method described in section 4.4.2. It is demonstrated that an improvement of 15 – 20% in LY can be achieved with DDA at concentrations between 0.1 – 0.5 DDA:Te MR. The uncertainty is dominated by sample to sample variation most likely arising from weighing error during sample preparation.

This result was subsequently confirmed by an independent LY measurement at University of Pennsylvania of samples prepared at Laurentian University [92].

4.5.2 Stability Improvement with DDA

To test the stability of the scintillator cocktail, samples were exposed to humid conditions. Various amounts of DDA were added into identical samples of 1% Te LAB with 2 g/L PPO. The 3 g sample cuvettes were then opened, placed into a large beaker with ~ 1 cm of water, and covered with plastic wrap, where they were left for a predetermined amount of time. Typically a 48 hour long exposure was chosen in order to induce a crash in a non-protected (without any amine) loaded sample within a few weeks.

The humidity exposed samples were kept at a constant temperature of 10°C and monitored visually over the course of 3 years. Furthermore, it is shown below that LY change over time can be used as a diagnostics of sample stability.

4.5.2.1 Visual Inspection

The increased tolerance to humidity exposure of Te-loaded scintillator with DDA has been tested in tens of samples and is illustrated in Figure 4.13, which shows photos of the exposed samples over time. While vials without DDA develop visible signs of deterioration after less than 2 months following a 48 hours humidity exposure, there has been no flake formation observed in samples with DDA:Te MR > 0.25 after 3 years of testing. Therefore, it is concluded that the addition of DDA can dramatically increase the tolerance of Te-loaded scintillator to humidity exposure by at least a factor of 18 (samples with DDA stable for more than 36 months compared with 2 months for samples without DDA).

The tests were also repeated with different durations of humidity exposure and yielded consistent results. Furthermore, extensive tests using visual inspection were set up at Queen's University with material synthesised at Laurentian University and also confirmed the positive effect of DDA on humidity resistance of Te-loaded scintillator. To date, no flake formation was ever observed in DDA protected samples (DDA:Te MR ≥ 0.25) with tests lasting up to 3 years in a

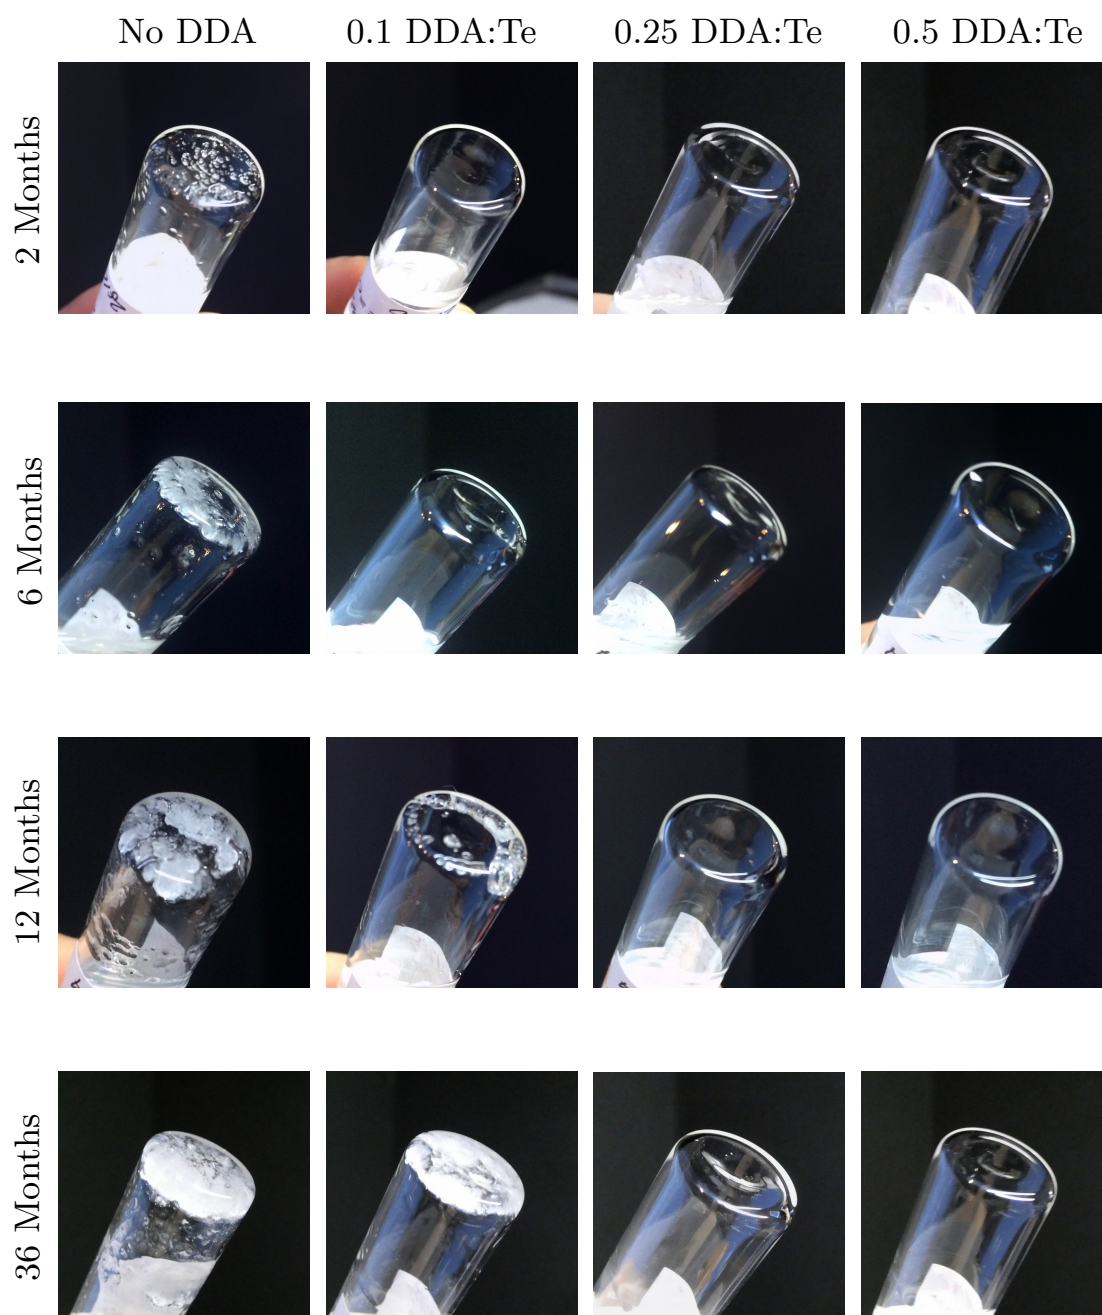


Figure 4.13: Photos of humidity exposed 1% Te samples with variable amounts of DDA. Each row shows the samples after a different amount of time following a 48 h humidity exposure. Significant flake formation is visible in the sample with no amine added (far left column). Sample with 0.1 DDA:Te MR (second column from left) developed flakes after ≤ 12 months and samples 0.25 (second column from right) and 0.5 (far right column) DDA:Te MR remain intact after 3 years. While only one set of samples is shown, equivalent behaviour is observed with duplicate samples. Furthermore, samples with higher concentrations of DDA than 0.5 MR also remain stable, but are not shown.

variety of environmental conditions.

4.5.2.2 LY Monitoring as a Measure of Stability

If no structural changes are present, the LY of a given sample is expected to remain constant over time. A gradual light yield increase over time was previously observed in some samples that eventually developed flakes. It is hypothesised that the flakes are formed from long Te polymers and that the LY increase is a result of an effective decrease of tellurium concentration in the solution and, thus, decreased quenching. It was indeed confirmed that the flakes have crystalline structure and high tellurium content [90].

In some cases, LY increase can be an early indicator of "phase separation" rather than flake formation. In this case, an abundance of water drives the hydrolysis reaction and results in immiscibility of the complex in LAB. Visually this manifests as an interface between the organic layer on top and aquatic on bottom. The tellurium concentration in the scintillator layer and, hence, the quenching, is decreased. Common to phase separation and flake formation, increase in LY is a sign of segregation of tellurium from the solvent and can be used as an early indication of instability. It is noted that in this case a gradual LY increase over time following a humidity exposure is discussed as a monitoring tool, which is different to the immediate LY increase after DDA addition described in section 4.5.1. For monitoring, the LY of the sample is measured repeatedly and compared to previous measurements for the same sample.

Figure 4.14 shows the light yield change of humidity exposed samples, which were periodically measured in addition to visual inspection. On the left, the light yield of these samples is shown in the arbitrary ADC units. It is demonstrated that while the samples with DDA have higher LY prior the exposure (time = 0 days), samples without any amine have steeper LY increase over time. This relative LY change of each sample is better seen on the plot on the right, where each

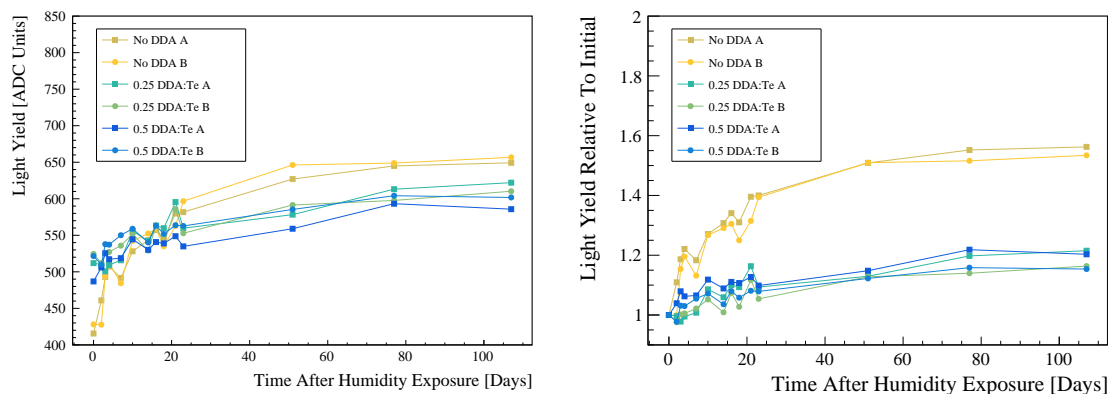


Figure 4.14: Light yield change after humidity exposure of 1% Te samples with no DDA (oranges), 0.25 (greens) and 0.5 (blues) DDA:Te MR. Both absolute LY in ADC units (left) and LY at the given time divided by an initial LY measurement of the sample prior to humidity exposure (right) are shown. Duplicate samples (A and B) are shown for error estimation.

measurement was divided by the initial LY of the given sample. It is observed that samples without any amine have significantly greater LY change indicating instability compared to samples with DDA as early as 3 days after humidity exposure.

4.5.3 Mechanism

The mechanism by which the amine protects the scintillator cocktail from water exposure and ensures higher LY is a subject of ongoing studies. However, it was confirmed by NMR that DDA forms an ionic association with the TeBD in water [90]. In LAB, studies hint towards hydrogen bonding between DDA and TeBD. The tellurium compounds are acidic and, hence, the OH groups are deprotonated. On the other hand, the amine is a base with a protonated site at the nitrogen atom. Hence, it is believed that the amine associates with the free OH groups of the TeBD and prevents further polymerisation (and thus crashing).

In the case of DDA, a long hydrocarbon chain is also associated to the complex, which further improves the solubility in the organic solvent. Moreover, it is possible that the presence of the tail prevents hydrolysis by sterically blocking

water molecules from accessing the bonds between Te and BD.

The reduction of scintillator light quenching with DDA can also be attributed to the neutralisation of the TeBD acidity. It is hypothesised that the distortion of π electrons of LAB aromatic rings decreases as the TeBD complex is neutralised. Alternatively, the decrease in quenching can be a result of the aforementioned steric blocking, reducing the interaction between the TeBD complex and the scintillator molecules.

4.6 Underground Processing

Great care has to be taken in order to ensure high radiopurity of all the components of the Te-loaded scintillator. In particular, cosmic ray spallation on tellurium can lead to the production of daughter nuclei with long half-lives and decay energies close to the Q-value of ^{130}Te . For this reason, the tellurium has been shipped underground in the form of telluric acid and it has been stored there since 2015. During this "cool down" period in a low cosmic ray environment, the nuclei activated on surface have a chance to decay before the scintillator loading with tellurium commences.

Two chemical processing plants have been built in the SNOLAB underground laboratory for the purpose of tellurium treatment: the TeA purification plant and the TeBD synthesis plant. The former will be used to remove K, U and Th chain impurities as well as remaining isotopes from cosmic ray spallation. The purification process consists of filtration followed by two recrystallisations (thermal and acidic) to remove both water soluble and insoluble contaminants. The procedure is described fully in [96].

The aqueous solution of the purified TeA will subsequently be transported to the TeBD plant, where it will be used for synthesis as described in section 4.4. 1,2-Butanediol, purified in the underground system used for LAB distillation, will be combined with TeA in the reaction chamber of the TeBD plant. A

recirculation loop will be responsible for passing the mixture through a heater and water vapour from the reaction will be collected in the condenser above the chamber. Once synthesised, the TeBD complex will be diluted at a 1:1 ratio with the scintillator mixture from the AV and then returned back to the detector. DDA can be added directly to the AV either before or after the loading process.

4.7 Summary

The composition and synthesis of Te-loaded scintillator for $0\nu\beta\beta$ search in SNO+ was described in this chapter. It was reported that by adding DDA, the LY of the tellurium loaded scintillator mixture can be increased by 15 – 20% with respect to samples with the same concentration of tellurium and no DDA. It follows that introducing this amine is expected to result in improved sensitivity to the $0\nu\beta\beta$ half-life for a given loading amount. Furthermore, it was shown that long-term stability of the loaded scintillator is improved by more than an order of magnitude when combined with the same ratio of DDA. Therefore, the inclusion of DDA into the scintillator cocktail is highly desirable and the SNO+ collaboration now plans its addition at a concentration of 0.5 DDA:Te MR.

In the next chapters, a Te-loaded scintillator mixture consisting of LAB + 2 g/L PPO + 0.5 % Te + 15 mg/L bisMSB + 0.5 MR DDA:Te with light yield of 480 PMT hits/MeV is assumed unless stated otherwise.

Chapter 5

Events in SNO+

Once the detector is filled with Te-loaded scintillator, the search for $0\nu\beta\beta$ will start. In order to apply statistical analyses on the collected data, the primary particle interactions inside the scintillator have to be understood, along with how to translate the PMT hits recorded in a given event into physical variables, such as the initial energy of the particle.

This chapter first lists the different interactions that can trigger the detector and lead to the creation of an event. Next, the algorithms responsible for reconstructing the event's energy and position inside the detector are described. The performance of these tools is outlined as they are a source of uncertainty for the $0\nu\beta\beta$ search.

Finally, event classification techniques can be employed to distinguish between the various event types using the differences in their timing profile, topology, coincidence with other events or in the path lengths of the primary particles responsible for the events. The last parts of this chapter focus on the implementation of two such techniques to help discriminate between $0\nu\beta\beta$ events and radioactive decays involving γ emission.

5.1 Types of Events in SNO+

The passage of particles through the scintillator inside the acrylic vessel of SNO+ is detected via optical photons emitted as scintillation or Čerenkov radiation. However, the number of detected Čerenkov photons is $< 3\%$ of the number of detected scintillation photons for a given event at energies close to the ^{130}Te Q-value. The event generating this light will typically belong to one of the categories below.

Electrons

Electrons are produced in single β decays of radioactive backgrounds and neutrino interactions as well as $2\nu\beta\beta$ and $0\nu\beta\beta$ processes. They travel distances of ~ 0.5 cm/MeV in liquid scintillator and emit scintillation light according to equation 4.1.

α Particles

Like electrons, α particles scintillate as they propagate through the scintillator, but their energy deposit per unit length is larger. Higher ionisation density leads to smaller light output than expected, which is parametrised by a larger Birk's constant in equation 4.1. As a result, for a given kinetic energy, the light output of an α particle is reduced compared to an electron. Furthermore, the ionisation density also affects the scintillator emission time as described in section 4.1. This creates a potential for pulse shape discrimination between α and β particles [97].

γ s

Uncharged particles do not scintillate directly. However, γ radiation produces secondary electrons as it undergoes Compton scattering and pair production in the scintillator. In contrast to electrons and α s, the average path length of a

1 MeV γ is ~ 30 cm, along which it would typically undergo multiple Compton scatters.

Positrons

Positrons scintillate as they travel through the scintillator inside the AV in the same way as electrons. However, at the end of their trajectory, they annihilate with surrounding electrons, producing two γ s. Consequently, the energy deposited in the scintillator is ~ 1 MeV larger than in the case of an electron of the same kinetic energy.

Neutrons

Neutrons in SNO+ are produced in inverse double beta decay ($\bar{\nu}_e + p \rightarrow e^+ + n$), hadronic showers after a muon passage, and (α, n) reactions, where an α interacts with a nucleus, leading to a neutron emission. Free neutrons scatter off protons along their trajectory, causing the protons to scintillate as they recoil. Once slowed down, the neutron captures on a hydrogen atom in $> 99\%$ cases, which then de-excites via a 2.2 MeV γ emission [60].

Muons

In the deep location of SNOLAB, the expected rate of cosmic muons passing through the central 8.3 m of the detector is 2.6 per hour [98]. Muons deposit energy along their path through the detector, producing large number of photons and triggering the detector multiple, consecutive times. Furthermore, compared to events contained inside the detector volume, muons entering from and exiting to the outside of the PSUP trigger the outward looking PMTs. As a result, events involving muons can be easily recognised and rejected. Hence, these events are not considered further in this work.

Radioactive Decays

In many cases, there are events that involve more than one primary particle. In particular, radioactive decays are of interest as they act as a background to rare event searches. Often, the initial decay particle, which converts the nucleus into a daughter isotope, is followed by emission of one or more γ s as the daughter nucleus de-excites. As an example, the radioactive isotope ^{208}Tl undergoes a β^- decay to ^{208}Pb . Depending on the excited state of the daughter nucleus, this is immediately followed by a cascade of 2-5 γ s. Similarly, ^{60}Co isotope decays to ^{60}Ni via $\beta\gamma\gamma$ emission in 99.8% of cases. The decay schemes of these two isotopes are shown in Appendix C. Importantly, the emission of the γ s makes it possible to distinguish these decays from the signal, which only involves electrons, as is further discussed in section 5.3. The full list of radioactive isotopes relevant for $0\nu\beta\beta$ search is presented in section 6.1.

Instrumental Events

Additionally, the detector can be triggered by events associated with the PMTs and electronics, rather than the passage of a primary particle through the detector itself. These events are referred to as "instrumental". A common example is a PMT "flasher", a static discharge inside the PMT leading to the emission of light into the detector. Another frequent source of instrumental events is electronic pick up, usually resulting from disturbances due to human activity near the crates. Instrumental events can be identified from the time, charge and spatial distributions of the PMT hits. Consequently, these events are removed from the data prior to analysis by a set of algorithms called data-cleaning, which only relies on the low-level PMT information (PMT time, charge and location).

5.2 Event Reconstruction in Liquid Scintillator

In order to convert PMT hits into meaningful physical quantities, such as energy and position of the original interaction that caused the detector to trigger, reconstruction algorithms have been developed. The reconstruction in SNO+ uses Cartesian coordinate system with origin in the centre of the PSUP. The XY plane is horizontal, while the positive Z axis points up through the neck. Due to the mostly spherical geometry of the detector, the radial distance from the origin, $R = \sqrt{X^2 + Y^2 + Z^2}$, is also often used. Additionally, polar angle θ , measured from the Z axis, is occasionally employed. This is particularly useful for situations where the use of spherical geometry is practical and where it is necessary to account for the asymmetry in Z due to the presence of the neck.

The SNO+ collaboration developed a set of algorithms incorporated into RAT called `ScintFitter` for energy, position and time reconstruction of events in both unloaded and Te-loaded scintillator. Due to the isotropic emission of scintillator light, observed signals are not sensitive to the directions of the original particles. All events are reconstructed under the assumption that the particle was an electron. This does not typically impact the interpretation of event's position, but it can change the interpretation of the reconstructed energy. Therefore, it is important to note that this reconstructed parameter is expressed in terms of the equivalent electron energy. The two main components of `ScintFitter` are described below.

5.2.1 Position and Time Reconstruction

The position and time of an interaction vertex are reconstructed first, as they are required as input to the energy reconstruction algorithm. For this purpose, a maximum likelihood fitter, called `PositionTimeLikelihood`, is employed [77].

For each event, the PMT hits are converted into time residuals, defined as

follows:

$$t_{res} = t_{hit} - t_{fit} - t_{tof}, \quad (5.1)$$

where t_{hit} is the time when the PMT was hit as recorded by the electronics and t_{fit} is the reconstructed time of the underlying interaction, which resulted in the creation of photons. While t_{hit} and t_{res} are different for each triggered PMT, t_{fit} has a global value for all the hits as it reflects the event time. t_{tof} is the time of flight of the photon from the reconstructed event vertex to the PMT and acts as a correction to t_{hit} for the distance the photon must have travelled to reach the particular PMT. Straight line paths between the vertex and PMTs are assumed and the variable photon speed in the different materials (scintillator, acrylic and water) is taken into account when calculating the time of flight. With perfect reconstruction, detector response, and no re-emission and scattering, the time residuals would give the emission time profile of a point-like event.

Time residuals are a useful concept utilised throughout the experiment. In the context of `PositionTimeLikelihood`, a hypothesised vertex (in space and time) is first proposed in the given fitter step and then used to calculate t_{tof} and t_{res} . The log-likelihood is then calculated as:

$$\log \mathcal{L} = \sum_{i=0}^{N_{Hits}} \log P(t_{res}^i), \quad (5.2)$$

where $P(t_{res}^i)$ is the probability of a hit having time residual t_{res}^i and N_{Hits} is the number of PMT hits in the event. The probability is taken from a probability distribution function (PDF) that is saved in the RAT database. This time residual PDF is the combination of the scintillator timing profile and detector effects. It is obtained from Monte Carlo simulation by calculating the time residuals using the true position and time of the event in the time residual calculation. The fit result of `PositionTimeLikelihood` maximises the agreement of the time residuals between the reconstructed vertex and the PDF using the Powell optimisation

method [99].

The performance of this fitter is demonstrated in Figure 5.1, which was obtained from comparison between true and reconstructed positions of events simulated with RAT version 6.16.3. To obtain the performance as a function of the radial distance of the interaction, 2.5 MeV electrons were simulated in 15 spherical shells of the same thickness. Conversely, for the energy dependence of the reconstruction quality, 1 – 10 MeV electrons were simulated homogeneously in the AV within a 5.5 m radius. For each radial shell, energy, and position coordinate (X,Y or Z), a histogram of reconstructed minus true position was made and fitted with a Gaussian distribution using MINUIT in ROOT. The mean of the Gaussian is defined as the bias and the standard deviation as the resolution.

Resolutions of 6 – 8 cm are observed in all 3 directions, resulting in overall position resolution of 10 – 14 cm. The best results are achieved at ~ 3 MeV, because the time residual PDF used by the fitter is produced using electrons of that energy. This choice was made so that the optimum lies close to the ^{130}Te Q-value. The main requirement on position resolution comes from its impact on energy resolution. The results in the next section demonstrate that this requirement is satisfied with the described position reconstruction. The mean position bias is of the order of mm in X and Y and 2 cm in Z due to the asymmetry of the detector in this coordinate. In all cases, this is much smaller compared to the resolution and, therefore, not problematic. The bias in Z can also be decreased by simply applying a correction to it, but this has not been explored here.

5.2.2 Energy Reconstruction

Energy is the primary observable for all of the physics searches in liquid scintillator in SNO+ and, therefore, event energy reconstruction is a crucial step in data analysis. The algorithm responsible for energy reconstruction in `ScintFitter`, `EnergyRThetaFunctional` [100], is described below.

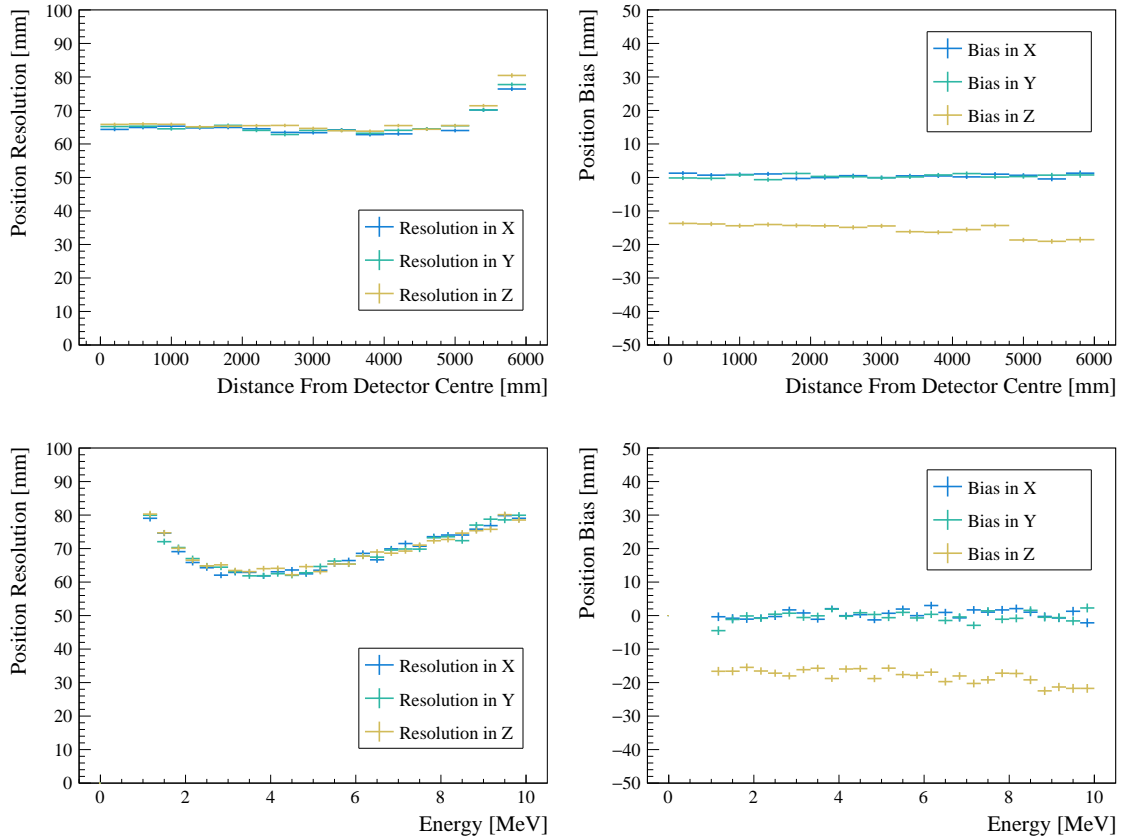


Figure 5.1: Position resolution and bias of electron event reconstruction in Te-loaded scintillator as a function of electron’s radial position (top) and energy (bottom). Each point represents the result of a Gaussian fit to a histogram of the difference between true and reconstructed position. Vertical error bars on the bias and resolution arise from errors of the Gaussian fit for the mean and the standard deviation, respectively.

For the liquid scintillator mixture in SNO+, the Birk’s constant introduced in equation 4.1 is of the order 10^{-4} m/MeV for both electrons and α particles [57]. Therefore, the relation between energy deposited in the scintillator and the number of photoelectrons created in the PMTs is very close to linear for energies of a few MeV. Unfortunately, the SNO PMT charge information is insufficient to accurately establish how many photoelectrons were created in a PMT and caused it to trigger. For most events, the number of triggered PMTs in the event, N_{Hits} , is relied upon to indicate the number of photoelectrons.

The relation between N_{Hits} and energy of simulated electrons is shown in Figure 5.2. The deviation from linearity is caused by hits with more than one

photon incident on a PMT. To overcome this, `EnergyRThetaFunctional` uses a multi-hit correction to obtain a linear energy estimator, H , as:

$$H = - \sum_{i=1}^S N_{PMT}^i \log\left(1 - \frac{N_{Hits}^i}{N_{PMT}^i}\right), \quad (5.3)$$

where the detector is divided into S segments containing PMTs with approximately equal solid angle with respect to the event position [101]. N_{Hits}^i are the number of triggered PMTs in segment i and N_{PMT}^i is the total number of PMTs in that segment. The multi-hit correction uses the fact that the number of photons hitting a PMT follows a Poisson distribution centred around a mean common to all PMTs in a segment. A full description can be found in [101]. The resulting estimator has a linear dependence on energy, as shown in Figure 5.2. The constant of proportionality between energy and H in the centre can be obtained from Monte Carlo simulation and verified using calibration sources.

The average number of photoelectrons also depends on the position of the event vertex due to variation in solid angle and photocathode coverage at the top of the detector. Moreover, absorption, scattering and total reflection probabilities depend on the photon path length. Therefore, to extend the reconstruction to other positions in the detector other than the centre, a correction is applied to translate the observed H at the vertex position to the one in the centre. A piecewise (inside and outside the AV) functional form in R is used in combination with a look up table in the polar angle θ . Again, the parameters of the functional form and table entries are determined from simulation.

Additionally, the number of operational PMTs and their thresholds are constantly fluctuating. As the fitter parameters are obtained from Monte Carlo with a fixed status of each PMT, the fit result has to be corrected for the variation in this setting. Appendix D describes the implementation of this correction and its validity.

The resulting performance of energy reconstruction with `ScintFitter` is




Figure 5.2: Number of triggered PMTs as a function of energy (blue) for electrons simulated at the centre of the detector. At higher energies, the probability of hits involving more than one photoelectron increases and, hence, the departure from linear relationship is more pronounced. A linear energy estimator H shown in orange.

summarised in Figure 5.3. Fractional resolution and bias are shown, defined as standard deviation and mean of the $(E_{reconstructed} - E_{true})/E_{true}$ distribution, respectively. Similarly to section 5.2.1, the dependence on R was obtained with 2.5 MeV electrons, whereas a uniform distribution of 1 – 10 MeV electrons filling a 5.5 m radius was used to show energy dependence.

The energy resolution is $\sim 3\%$ for energies around ^{130}Te Q-value. Given the light yield of the Te-loaded scintillator described in the previous chapter, 480 PMT hits/MeV are expected. Therefore, the Poisson limit on the resolution at the $0\nu\beta\beta$ energy is

$$1/\sqrt{N_{\text{Hits}}} = 1/\sqrt{2.5 \text{ MeV} \times 480 \text{ MeV}^{-1}} = 2.9\%. \quad (5.4)$$

It is hence concluded that the energy resolution of ScintFitter is very close to

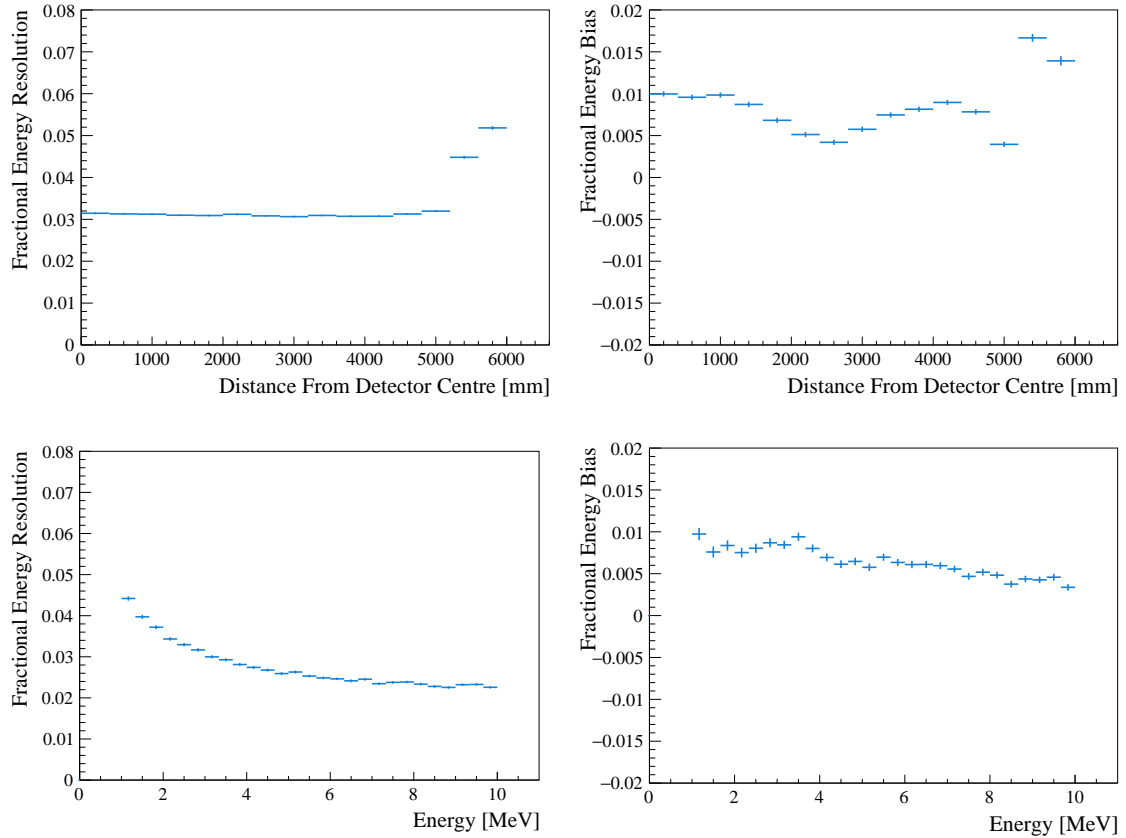


Figure 5.3: Energy resolution and bias of electron event reconstruction in Te-loaded scintillator as a function of electron’s radial position (top) and energy (bottom). Each point represents the result of a Gaussian fit to a histogram of $(E_{reconstructed} - E_{true})/E_{true}$, where $E_{reconstructed}$ and E_{true} are the reconstructed and simulated energy, respectively. Vertical error bars on the bias and resolution arise from errors of the Gaussian fit for the mean and the standard deviation, respectively.

the achievable limit. The energy bias is smaller than the resolution across the tested range of energies and positions.

5.3 Event Classification

Once the position and energy of an event are reconstructed, additional algorithms can be used to distinguish between different event types. In particular, the differences between primary particle behaviour inside the detector described in section 5.1 can be used to distinguish between them. In this section, the

focus is on distinguishing between events with point-like and multi-site energy deposition. This is of particular interest for distinguishing between $0\nu\beta\beta$ electronic events and radioactive background events involving γ rays. For this work, two classifiers (jointly referred to as the multi-site classifiers) were added into the RAT framework based on the work in [102] on this topic and are described below.

5.3.1 Timing Classifier

As the distance travelled by a single electron is shorter than the position resolution, its energy deposition can be considered point-like. The same is true for two electrons created simultaneously such as in the case of double β decay events. On the other hand, this description is not accurate for many radioactive backgrounds as they typically involve γ s with $\mathcal{O}(1 \text{ MeV})$ energies. In these events, there would typically be an energy deposit at the place where the decay happened from the β^- , β^+ or α . Additionally, energy will be deposited along the trajectory of the γ (s) coming from the multiple Compton scatters.

The position reconstruction, however, treats every event as an electron and only one position vertex is output for each event. The fitted position can then be used to calculate time residuals as defined in equation 5.1, again assuming that all the photons originated at a single vertex. If this assumption is not valid and the energy deposition in the scintillator is more spread out, the t_{res} distribution will appear wider than expected. This effect of time residual broadening is the basis for classifying these events.

Two radioactive backgrounds that are particularly important for the $0\nu\beta\beta$ search are considered here to demonstrate the technique: ^{60}Co decays from within the scintillator volume and "external" ^{208}Tl from outside of the scintillator volume. In the former decay, two γ s of 1.2 and 1.3 MeV are produced 99.8% of time in addition to a β^- with Q-value of 0.3 MeV (decay scheme in Figure C.1




Figure 5.4: Probability distribution of reconstructed energy for ^{60}Co (green) and $0\nu\beta\beta$ (orange, dashed) events in the scintillator and ^{208}Tl (blues) decays originating outside of the scintillator: in the acrylic vessel (AV) wall, in the hold-down ropes (HDR) and the external water. Only events with reconstructed distance < 5.5 m from centre of the AV are plotted. Distributions were created with 10^4 simulated events and error bars reflect the statistical fluctuations.

of Appendix C). The mean overall reconstructed energy is only ~ 30 keV away from the $0\nu\beta\beta$ peak at 2.527 MeV. The situation is different in the case of ^{208}Tl , where the Q-value of the decay is 5.0 MeV. Therefore, when these decays happen in the scintillator volume, a broad energy spectrum is detected, which can pose a statistical background to the search, but not mimic the signal peak. However, if ^{208}Tl decays outside of the scintillator volume, there is a chance of some of the γ s from the decay to penetrate into the scintillator volume and deposit their energy there. In particular, there is a 2.6 MeV γ produced in 99.8% of decays (decay scheme in Figure C.2 of Appendix C). Given the energy resolution of $\sim 3\%$, if this γ reconstructs inside the scintillator it can appear as the $0\nu\beta\beta$ signal.

The reconstructed energy of these two event types are compared to $0\nu\beta\beta$ in Figure 5.4. The ^{208}Tl events were simulated in three places outside the scintillator

volume: inside the AV wall, in the water surrounding the AV and inside the hold-down ropes (HDR). Only the small proportion of events that reconstructed within the inner 5.5 m radius of the AV is plotted. It is clear that both of these backgrounds can be mistaken for $0\nu\beta\beta$ if the search is only based on energy and, hence, the ability to discriminate between these events by other means is of high importance to the experiment.

Figure 5.5 shows the time residuals of simulated 2.5 MeV e^- , $0\nu\beta\beta$ and ^{60}Co events, calculated using their reconstructed position. Hits with time residuals below or above the axis range, were placed into the first (underflow) and last (overflow) bins, respectively. The broader peak of the ^{60}Co decays compared to e^- interactions is caused by the fact that the former results in multi-site energy deposition. It is also shown in the plot that there is no significant difference between $0\nu\beta\beta$ events and single electrons of the same energy. Only events with reconstructed radius < 4.2 m and with energies between 2.438 and 2.602 MeV are included in this plot. This is region surrounds the expected $0\nu\beta\beta$ peak in energy and contains a larger detector volume than what is typically considered in a counting analysis (3.3 or 3.5 m), where additional rejection of external backgrounds is expected to have the biggest impact on the sensitivity to $0\nu\beta\beta$. For Figure 5.5, the number of simulated events passing the selection is $\sim 10^4$ for each event type. The same selection of events is used throughout the remainder of this chapter.

A similar effect is seen with external ^{208}Tl events reconstructing inside the given radius, as shown in Figure 5.6. In addition to the time residual peak broadening, there is also a more prominent excess of PMT hits in the underflow bin for the ^{208}Tl events corresponding to early hit PMTs. In both ^{208}Tl and ^{60}Co events, the early hits are attributed to light coming from the decay β particle or early Compton scatters. However, in the case of ^{208}Tl , the selection only includes events that travelled far enough to reach the inner parts of the detector and,




Figure 5.5: Time residual spectra for ^{60}Co decay (blue), $0\nu\beta\beta$ (orange, dashed) and 2.5 MeV e^- events (yellow, full) inside 4.2 m radius. First and last bins are underflow and overflow, respectively. Distributions obtained from 10^4 MC events, but are normalised to one event. Error bars arising from statistical fluctuations are smaller than the line width of the histograms and, hence, are not shown.

hence, reconstructed far away from the initial decay (in both time and space). As a result, the light from the initial decay appears even earlier relative to the main peak than in the case of ^{60}Co . Even though $\mathcal{O}(10^8)$ of external ^{208}Tl events were simulated, only ~ 300 events of each external source reconstructed close enough to the detector centre to pass the event selection and be included in Figure 5.6.

The timing classifier calculates the time residuals for each event using the reconstructed position and uses them to evaluate how point-like or multi-site the event appears. This has been implemented as a linear classifier, which assigns a discriminant value to each event. The discriminant is calculated as

$$D = \vec{w} \cdot \vec{T}, \quad (5.5)$$

where \vec{T} is the vector of bin contents from the time residual histogram for the




Figure 5.6: Time residual spectrum for $0\nu\beta\beta$ events (orange) compared to ^{208}Tl backgrounds from the AV wall (purple), hold-down ropes (blue) and water surrounding the AV (green). Only events that reconstructed within 4.2 m radius and with energy between 2.438 and 2.602 MeV are shown. First and last bins are underflow and overflow, respectively. Distributions obtained from multiple MC events (10^4 for $0\nu\beta\beta$ and ~ 300 for each ^{208}Tl source), but are normalised to one event. Error bars arising from statistical fluctuations are smaller than the line width of the histograms and, hence, are not shown.

event to be classified and \vec{w} are weights that can be obtained from Monte Carlo simulation. The framework allows freedom in the choice of weights and two types are considered for this work. A simple choice is

$$w^i = \log \left(\frac{\mu_{0\nu}^i}{\mu_{bkg}^i} \right), \quad (5.6)$$

with $\mu_{0\nu, bkg}^i$ being the probability of a PMT hit in time residual bin i for a signal and background, respectively. $\vec{\mu}_{sig, bkg}$ are therefore simply vectors of bin contents obtained from time residual PDFs such as those in Figures 5.5 and 5.6. This way,

the discriminant becomes simply the log-likelihood ratio:

$$D = \Delta \log \mathcal{L} = \sum_{i=0}^{N_{Hits}} \log \frac{\mu_{0\nu}^i}{\mu_{bkg}^i}. \quad (5.7)$$

According to the Neyman-Pearson lemma, the log-likelihood discriminant is the most powerful technique under the assumption that the two hypotheses, signal and background, are described completely by the PDFs [103]. However, especially for the multi-site events, it is not guaranteed that the average PDF provides a good characterisation of individual events. For example, some multi-site events travel further than others and as a result have more hits in the later bins and fewer in the peak. Therefore, there is a correlation between the bins which is ignored by the log-likelihood. An alternative is to use the Fisher discriminant [104], which takes into account these correlations by including the covariance matrix in the calculation. In this case the weights are calculated as

$$\vec{w} = (\Sigma_{bkg} + \Sigma_{0\nu})^{-1} \cdot (\vec{\mu}_{0\nu} - \vec{\mu}_{bkg}), \quad (5.8)$$

where $\vec{\mu}_{0\nu, bkg}$ are again the vectors of mean PMT hit probability from the PDFs and $\Sigma_{0\nu, bkg}$ are the covariance matrices with dimensions equal to the number of bins in the PDFs. For both signal and backgrounds, the elements of those matrices give the covariance between the bins in the time residual spectra and are calculated as

$$\Sigma^{ij} = \frac{1}{N_{Events} - 1} \sum_{n=0}^{N_{Events}} (T_n^i - \mu^i) \cdot (T_n^j - \mu^j), \quad (5.9)$$

where N_{Events} is the number of events of the relevant type (signal or background) and T_n^i is the content of the i^{th} bin of the individual time residual spectrum of an event n . As before, μ^i is the content of bin i in the mean time residual spectrum for the given event sample.

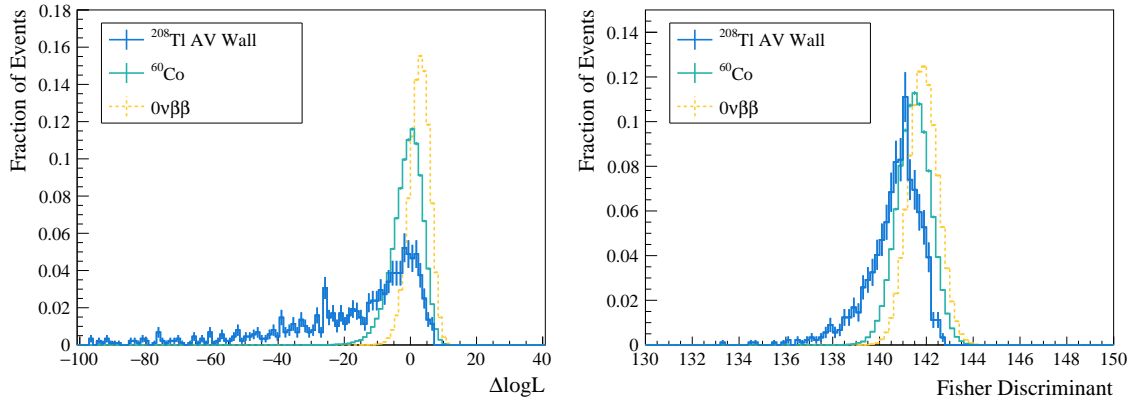


Figure 5.7: Discriminant values for $0\nu\beta\beta$ (orange), ^{60}Co (green) and external ^{208}Tl (blue) events. $\Delta \log \mathcal{L}$ is plotted on the left and Fisher discriminant on the right. Histograms have been normalised to unity and give the probability of an event of the specific type being assigned the given discriminant value. Poisson errors are shown.

Both log-likelihood and Fisher weights were obtained using Monte Carlo generated external ^{208}Tl decays. Multiple sets, each obtained using events simulated in different part of the detector, were made available in the RAT database from where they could be simply accessed by the classifier. The timing classifier was then run over an independent sample of simulated $0\nu\beta\beta$, external ^{208}Tl and ^{60}Co events.

Figure 5.7 shows the discriminant values assigned to these events and demonstrates a clear difference between the three populations. For the results presented in this work, the classifier was set to use weights calculated from ^{208}Tl events in the AV wall. This choice is motivated by the fact that out of the ^{208}Tl external background sources, these events are expected to have the largest contribution to the ROI used in Chapter 6. The separation between ^{60}Co and $0\nu\beta\beta$ demonstrates that it is not necessary to use specific weights for each background in order to be able to discriminate between them to a reasonable extent. This discrimination largely comes from the shared features of multi-site event time residuals. The effect of using weights calculated with different types of multi-site events is explored in more detail in Appendix E.

Ideally, one would like to find a value of the discriminant such that, by rejecting events this way, all the background would be eliminated with no loss of signal. It is clear that this "event-by-event" classification is not possible as there is a non-zero overlap between the curves in Figure 5.7. However, it is possible to cut a significant proportion of the background with minimal loss in $0\nu\beta\beta$ events. Figure 5.8 shows the fraction of accepted background events as a function of signal efficiency, which was obtained by placing cuts at different values of D . The log-likelihood weights perform better than Fisher ones apart from the region with low signal efficiency for ^{208}Tl . It is hypothesised that the Fisher discriminant is not fulfilling its potential due to limited statistics ($\mathcal{O}(10^2)$ events) available for the creation of the background covariance matrix. With the log-likelihood discriminant, $\sim 15\%$ of ^{60}Co and $\sim 60\%$ of ^{208}Tl events that penetrate into the inner 4.2 m of the scintillator volume can be rejected with a signal loss of $< 1\%$. More importantly, the discriminant distributions allow for good statistical separation, which is discussed later in this work.

5.3.2 Topological Classifier

On top of the timing differences, the topology of background events coming from outside the AV volume is also different compared to $0\nu\beta\beta$ events inside the scintillator. The probability of the 2.6 MeV γ travelling the distance x through a material falls off approximately as e^{-kx} , where k is a constant. Therefore, the γ s with the largest chance of reaching the scintillator volume are those travelling directly towards the detector centre from the original decay, as depicted on Figure 5.9. Consequently, the early hits on PMTs in proximity to the original decay are most likely to be in the direction of a vector from the detector centre to the reconstructed vertex. Figure 5.10 shows the probability distribution of early PMT hits (time residuals between -2 and 20 ns) as a function of $\cos\theta$ for signal and external ^{208}Tl decays. θ is the angle between the hit PMT position vector and




Figure 5.8: Fraction of ^{60}Co (green circles) and ^{208}Tl events from the AV wall (blue squares) passing the timing classifier cut at a given discriminant value as a function of fraction of accepted signal events. Curves showed for $\Delta \log \mathcal{L}$ (dashed line) and Fisher (full line) weights, both tuned on an independent MC simulation of ^{208}Tl decays in the AV.

the reconstructed event position vector as marked in Figure 5.9. As expected, the background PMT hit distribution is strongly peaked towards $\cos \theta = 1$, whereas the signal distribution is approximately flat as the events are isotropic within the AV. The slight slope of the $0\nu\beta\beta$ distribution is believed to be due to optical effects. The topological classifier exploits the difference in the $\cos \theta$ distributions to help discriminate between signal and external backgrounds.

The implementation of the topological classifier is analogous to the timing case, with the main difference that the event's PMT hit distribution in $\cos \theta$, rather than time residuals, is used. In addition, there is now a cut that selects only early hits to be used in the classification. The discriminant is calculated as:

$$D = \vec{w}_\theta \cdot \vec{\Theta} = \sum_{i=0}^{N'_{Hits}} \log \left(\frac{\lambda_{0\nu}^i}{\lambda_{bkg}^i} \right), \quad (5.10)$$




Figure 5.9: A schematic of an external background event originating inside the AV wall. The two concentric black circles represent the AV and the PSUP and the small blue circles individual PMT hits. The reconstructed position of the event is given by the grey vector \vec{r} pointing away from the origin at the centre of the detector. For each PMT hit, θ is defined as the angle between \vec{r} and the PMT position vector (green). PMTs with $\cos \theta \simeq 1$ are most likely to register early hits.

where N'_{Hits} is the number of early hit PMTs, w_{θ}^i are the topological weights and $\vec{\Theta}$ is the bin contents of the hit distribution in $\cos \theta$ for the event to be classified. Although the implementation still allows the freedom to add different type of weights, only $\Delta \log \mathcal{L}$ weights are considered here. These are again calculated from the average probability, $\lambda_{0\nu, bkg}$, for signal and background in the corresponding $\cos \theta$ bin.

Figure 5.11 shows the values output by the classifier for a sample of $0\nu\beta\beta$ and external ^{208}Tl decays. As with the timing classifier, events with radius < 4.2 m and energy between 2.438 and 2.602 MeV are included and the weights were calculated taking ^{208}Tl events originating in the AV wall as background. The efficiency of the classifier in rejecting the background events is plotted as a function of signal acceptance in Figure 5.12. The graph also shows rejection




Figure 5.10: Angular distribution of early PMT hits for simulated $0\nu\beta\beta$ (orange) and external background events from the AV wall (purple), hold-down ropes (blue) and water surrounding the AV (green). The horizontal axis represents the cosine of angle between each early hit PMT position and the reconstructed event vertex as measured from the centre of the detector. Distributions obtained from multiple MC events (10^4 events for $0\nu\beta\beta$ and ~ 300 events for each ^{208}Tl source), but are normalised to unity. Error bars reflect the statistical fluctuations.

potential for ^{60}Co for reference. As expected for events originating inside the scintillator volume, ^{60}Co cannot be distinguished using the topological classifier as the events are isotropic. For external ^{208}Tl however, about 60% can be rejected with minimal signal loss using the topological classifier alone.

Inevitably, the topological classifier will cut many of the same events as the timing one. To assess whether there is a benefit of using both classifiers in combination, the two discriminant values were plotted simultaneously for each ^{208}Tl AV event in Figure 5.13. Not surprisingly, there is a correlation between the two discriminant values for ^{208}Tl , owing to the fact that external events detected further from their origin have wider time residuals and are more likely to have γ s travelling directly towards the detector centre. On the other hand, there is no correlation between the two discriminants for the signal events. Therefore, it




Figure 5.11: Discriminant values output by the topological classifier for $0\nu\beta\beta$ and ^{208}Tl events from AV wall. Histograms have been normalised to unity and give the probability of an event of the specific type being assigned the given discriminant value. Poisson errors are shown.

should be possible to use this correlation to further improve the discrimination between signal and external backgrounds compared with what can be achieved with either of the classifiers alone.

5.4 Summary

SNO+ data is split into events, each corresponding to one detector trigger. There is a number of processes, including mostly primary particle interactions and radioactive isotope decays, that can cause an event. For each event in Te-loaded liquid scintillator, energy and position is reconstructed under the assumption that the underlying particle was an electron. The algorithms responsible for this reconstruction have been described above.

In some cases, it is also possible to use the PMT hit information to distinguish the type of the event, i.e. the original interaction. For this work, two linear classi-




Figure 5.12: Fraction of ^{60}Co (green circles) and ^{208}Tl events from the AV wall (blue squares) passing the topological classifier cut at a given discriminant value as a function of fraction of accepted signal events. Curves showed for $\Delta \log \mathcal{L}$ weights tuned on an independent MC simulation of ^{208}Tl decays in the AV.

fiers were implemented into the SNO+ software, discriminating between multi-site and point-like events based on timing and topological parameters. This is particularly useful for rejection of radioactive backgrounds involving photons in the $0\nu\beta\beta$ analysis, where the signal can be considered as point-like. It has been demonstrated that about 60% of external and 15% of internal background events can be rejected with minimal signal loss with the use of a single classifier. The inclusion of the classifiers in a $0\nu\beta\beta$ analysis and the impact on sensitivity of SNO+ to the given process is explored in detail in the following chapter.




Figure 5.13: Contour plot of timing and topological discriminant values for $0\nu\beta\beta$ (orange) and external ^{208}Tl (blue). Contours give lines of equal probability of an event to have the given combination of discriminants and are showed for $\Delta\log\mathcal{L}$ weights tuned on an independent MC simulation of ^{208}Tl decays in the AV.

Chapter 6

Likelihood Analysis for Neutrinoless Double Beta Decay

After reconstruction and classification of events from the detector, the final step in obtaining a physics measurement from the tellurium phase of SNO+ will be the extraction of the $0\nu\beta\beta$ signal from the data. At the time of writing this work, the experiment has not commenced data taking with the tellurium isotope. Nevertheless, the signal extraction for $0\nu\beta\beta$ can be tested on Monte Carlo simulation and the expected sensitivity to this process can be obtained.

In this chapter, a likelihood analysis for $0\nu\beta\beta$ is presented. First, the various backgrounds to this search are discussed, which provides an important input for the analysis design. Next, the signal extraction framework is described. Particular attention is paid to extending a two-dimensional analysis (in energy and radius) by including the classifier discriminants from the preceding chapter as additional observables. The analysis builds significantly on previous work that developed the two-dimensional signal extraction framework [101]. Finally, the

sensitivity of the experiment is estimated with regards to both setting a lower limit on $0\nu\beta\beta$ half-life and discovery of the process. The improvements achieved due to the inclusion of timing and topological observables are discussed in detail.

6.1 Backgrounds for $0\nu\beta\beta$

There are multiple sources of potential backgrounds for the $0\nu\beta\beta$ search in SNO+, all of which can be divided into two categories. Internal backgrounds are non-signal events originating inside the AV volume. Conversely, external backgrounds originate outside of the AV or inside the acrylic itself. This section details the different sources of both these background categories and possible mitigation strategies.

6.1.1 $2\nu\beta\beta$

The introduction of natural tellurium into the liquid scintillator will inevitably lead to the detection of a number of $2\nu\beta\beta$ events from ^{130}Te nuclei. As the two anti-neutrinos do not leave any signal in the detector, this decay looks exactly like the desired signal; the only difference is the energy deposited in the scintillator by the two daughter electrons. Due to finite energy resolution, however, a portion of these events will be smeared into the ROI.

The current best measurement of the $2\nu\beta\beta$ half-life in ^{130}Te comes from the CUORE experiment [41]. Based on this value of $T_{1/2}^{2\nu} = 7.9 \pm 0.1(\text{stat.}) \pm 0.2(\text{sys.})$ years, 5.4×10^6 $2\nu\beta\beta$ events per year are expected in the SNO+ detector containing 0.5% natural tellurium by weight. Even though only a small number of these events will reconstruct with energy inside the ROI, the $2\nu\beta\beta$ is still a major contribution to the overall background rate.

The sharp falling edge of the $2\nu\beta\beta$ spectrum close to the $Q_{\beta\beta}$ value means that

this background quickly dominates at the lower edge of the signal region. As a result, counting analyses that simply compare the number of detected events to the expected background rate typically set a narrow and asymmetric energy ROI around the expected $0\nu\beta\beta$ peak. However, in case of spectral analyses, where shapes of individual backgrounds are fitted based on Monte Carlo predictions, it is beneficial to widen the energy signal region. This way, the falling edge of the $2\nu\beta\beta$ spectrum actually helps constrain the number of background events of this type.

6.1.2 Solar Neutrino Scattering

Another important background comes from elastic scattering of solar neutrinos off electrons inside the scintillator material. These interactions produce a single recoil electron that scintillates and can act as a background for the $0\nu\beta\beta$ search. The two tree-level Feynman diagrams of the $\nu_x + e^- \rightarrow \nu_x + e^-$ process are shown in Figure 6.1. The neutral current interaction is available for all three neutrino flavours. On the other hand, the energies of solar neutrinos reaching the Earth are insufficient to produce a μ or τ charged lepton in the detector. Therefore, the charged current interaction is only available for electron neutrinos in the case of solar neutrino scattering.

In the laboratory frame, the differential cross-section of these diagrams is

$$\frac{d\sigma}{dT_e}(E_\nu, T_e) = \frac{2G_F^2 m_e}{\pi} \left[g_1^2 + g_2^2 \left(1 - \frac{T_e}{E_\nu}\right)^2 - g_1 g_2 \frac{m_e T_e}{E_\nu^2} \right], \quad (6.1)$$

where T_e is the kinetic energy of the recoil electron, σ is the interaction cross-section, G_F is the Fermi constant and m_e is the electron mass [8]. The values g_1 and g_2 are constants, but depend on the neutrino flavour, with larger values for ν_e than for $\nu_{\mu,\tau}$. Finally, E_ν is the incoming neutrino energy. Figure 6.2 shows the energy spectra of solar neutrinos reaching the Earth. The energy distribution of




Figure 6.1: Tree-level Feynman diagrams for neutrino elastic scattering on electrons. Charged current interaction (left) is only available to electron solar neutrinos due to low solar neutrino energies and lepton flavour conservation in the weak interaction. The neutral current interaction (right) is observed with all three neutrino flavours, i.e. $x = e^-, \mu^-, \tau^-$.




Figure 6.2: Flux of solar neutrinos upon reaching Earth as a function of energy. Data taken from [105].

the observed electron, T_e , is the convolution of the differential cross-section from equation 6.1 and the neutrino energy spectrum.

The interactions that contribute to $0\nu\beta\beta$ background will produce electrons with $T_e \sim 2.5$ MeV, which is only possible for neutrinos with energies above this value. As seen in Figure 6.1, the dominant contribution to the solar neutrino flux at such energies is from the ${}^8\text{B}$ neutrinos, which are produced in the Sun in the




Figure 6.3: The energy distribution of recoiled electrons from ${}^8\text{B}$ solar neutrino elastic scattering as output by RAT. Insert in top right corner zooms in on the energy region around ${}^{130}\text{Te}$ $Q_{\beta\beta}$.

following reaction:



The energy of the recoil electrons from ${}^8\text{B}$ neutrino elastic scattering is plotted in Figure 6.3. This distribution was obtained from RAT, which performs the convolution of the ${}^8\text{B}$ neutrino spectrum and the differential cross-section by Monte Carlo sampling. It is shown that around the ROI, the contribution to the background from these interactions is relatively flat in energy compared to other backgrounds. The total flux of ${}^8\text{B}$ neutrinos was measured by multiple experiments [10, 76, 106] and theoretical predictions are available from solar models. Based on the theoretical value from [9], around 1300 ${}^8\text{B}$ neutrino interactions are expected per year in the SNO+ AV volume across all energies [88]. However, the finite error on the flux measurements and predictions as well as the theoretical uncertainty on the shape of the ${}^8\text{B}$ spectrum results in some uncertainty on the exact number of events that are expected in the $0\nu\beta\beta$ ROI.

As discussed in section 5.3, the position resolution is insufficient to distin-

Source	Target Concentration [g/g]	Decays per Year
^{238}U chain	1.3×10^{-15}	39.52×10^4
^{232}Th chain	5.6×10^{-16}	5.57×10^4

Table 6.1: Expected internal U/Th background rates in Te-loaded scintillator cocktail. The assumed purity level of LAB is $\mathcal{O}(10^{-17}\text{g/g})$ as achieved by other liquid scintillator experiments [55, 109]. Level of $\mathcal{O}(10^{-13}\text{g/g})$ is expected for telluric acid after the purification procedure described in section 4.6.

guish between two electrons from $0\nu\beta\beta$ and a single electron of the same energy. Therefore, there is no observable difference in the scintillation light from the desired signal and ^8B neutrino scattering events in the ROI. One distinction of the solar neutrino scattering events is the Čerenkov light, which is emitted in a cone around the direction of travel of the electron. For ^8B events, the direction of the scattered electrons (and, hence, the Čerenkov light) is correlated with the Sun's position. However, there have been multiple studies deeming the discrimination between solar neutrino scattering and $0\nu\beta\beta$ events using directionality of Čerenkov light unfeasible in SNO+¹ [107, 108].

6.1.3 Internal U/Th Natural Radioactivity

The presence of natural radioisotopes inside the scintillator cocktail also results in events that can act as background for the $0\nu\beta\beta$ search. The main contributions come from daughters of the ^{238}U and ^{232}Th chains. The target levels of these contaminants in the Te-loaded scintillator are listed in Table 6.1. It is assumed that the decay chains are in equilibrium, i.e. the production rate of each daughter isotope is equal its decay rate. Therefore, the same number of events per year is expected for all the isotopes belonging to the same chain.

Figure 6.4 shows the ^{238}U and ^{232}Th decay chains, with the most worrisome isotopes highlighted in orange. The difference in timing between α and β scintil-

¹with the scintillator composition that will be used for the initial tellurium phase

lator excitations described in section 4.1 means that it is possible to reject events involving α particles. Therefore, only natural radioisotopes that undergo β decay (denoted by diagonal arrows in the figure) are of major concern. In particular, the $0\nu\beta\beta$ ROI is dominated by ^{212}Bi and ^{214}Bi β decays with Q-values of 2.25 and 3.27 MeV, respectively. Luckily, both of these decays result in production of short-lived polonium isotopes. ^{212}Po and ^{214}Po α decay with released energies of 8.95 and 7.83 MeV, which are quenched down to ~ 0.9 and ~ 0.8 MeV equivalent electron energies, respectively. The half-lives of the ^{212}Po and ^{214}Po isotopes are $0.3 \mu\text{s}$ and $164 \mu\text{s}$ respectively, and therefore, the α decay follows very shortly after the bismuth decay in both cases. Such short delay between two decays is uncommon and results in an easily recognisable signature in the detector. As a result, this coincidence between the polonium and bismuth decays can be used to identify and eliminate these backgrounds. Additionally, the rate of these coincidences provide a valuable measure of the contamination of the scintillator by the ^{238}U and ^{232}Th chain isotopes.

The $^{212/214}\text{Bi}$ and $^{212/214}\text{Po}$ decays can either be registered as two distinct events ("out-of-window") or fall into the same event window ("in-window"). In the former case, the two events will appear close to each other in both time and reconstructed position. A cut proposed in [110] rejects any pair of events where for the second trigger

$$N_{Hit} \geq 50 \quad \text{and} \quad \Delta t < 500 \text{ ns} \tag{6.3}$$

$$\text{or} \quad N_{Hit} \geq 50 \quad \text{and} \quad 500 \text{ ns} < \Delta t < 3.936 \text{ ms} \quad \text{and} \quad \Delta R < 1.5 \text{ m},$$

where Δt and ΔR are the differences in time and reconstructed position compared to the previous trigger. Reflected light produced in $0\nu\beta\beta$ events can sometimes cause the detector to trigger again, creating an event with low number of hit PMTs. The N_{Hit} condition in equation 6.3 ensures that signal events with re-triggers are not rejected. In the Monte Carlo simulation used throughout this

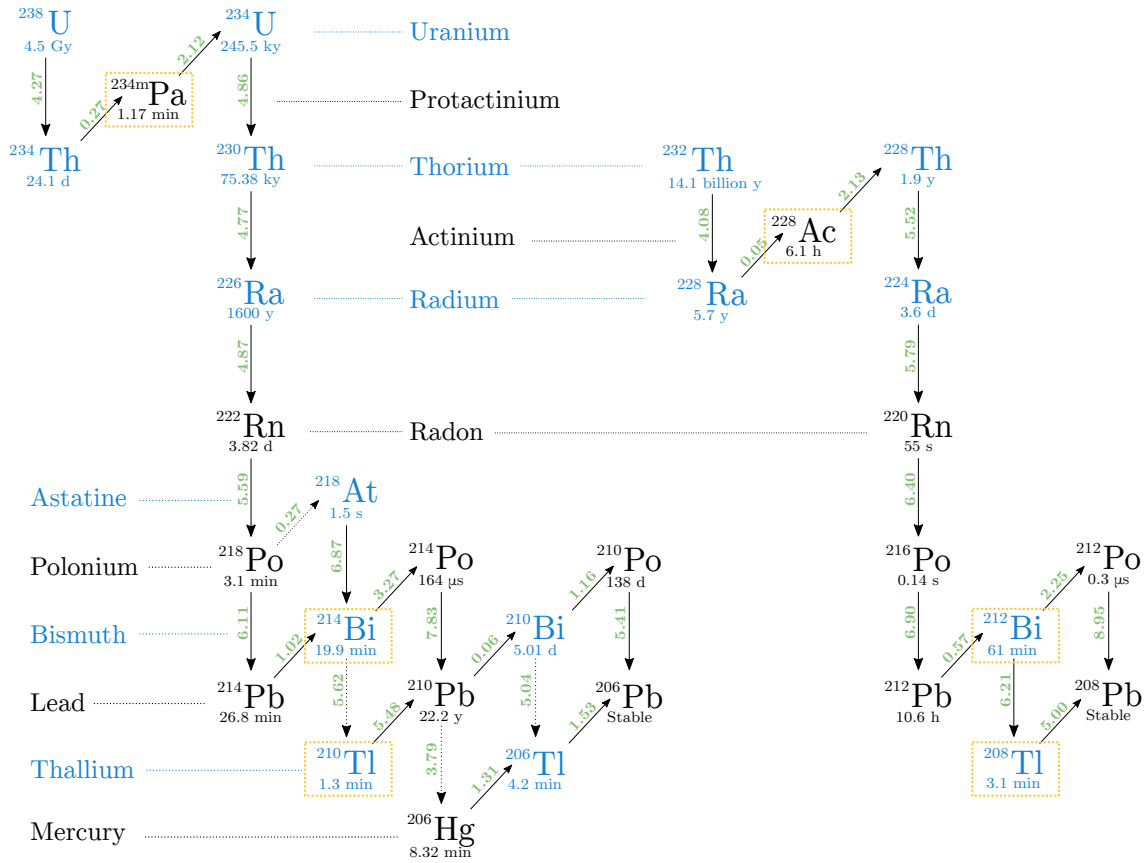


Figure 6.4: The ^{238}U (left) and ^{232}Th (right) radioactive decay chains. Vertical and diagonal arrows represent α and β^- decays respectively. Black numbers underneath isotopes give the half-lives and green numbers along the arrows give ground state to ground state Q-values of the decays in MeV. Decays of isotopes in orange boxes pose a potential background for the $0\nu\beta\beta$ search. Data taken from [64].

chapter, the cut above removes 6.4 % and 99.8 % of the out-of-window $^{212}\text{BiPo}$ and $^{214}\text{BiPo}$ events inside the ROI, respectively.

In the case where the bismuth and polonium decays are registered as single event, a double peak structure is seen in the event's time residuals. To reject these in-window coincidences, two classifiers, BiPoCumul and BiPoLikelihood214 [111], are used.

Figure 6.5 shows the energy distribution of the $^{212/214}\text{BiPo}$ decays and the effect of both in-window and out-of-window coincidence cuts. Jointly, these cuts are responsible for reduction of the $^{212/214}\text{BiPo}$ backgrounds in the ROI by 3-4 orders of magnitude. Due to the difference in ^{212}Po and ^{214}Po half-lives,




Figure 6.5: Reconstructed energy of $^{212/214}\text{Bi}$ β and $^{212/214}\text{Po}$ α (left/right) decays. Histograms obtained from RAT6.16.3, which generates the two decays in conjunction to simulate the coincidence. Only events reconstructed within a 5.5 m radius inside the AV and with valid ScintFitter result are included. Histograms are normalised to expected counts per year.

comparatively more $^{212}\text{BiPo}$ events are removed by the in-window classifiers, whereas the dominant reduction of $^{214}\text{BiPo}$ events come from the out-of-window cuts. It has been demonstrated that the signal loss from these cuts is negligible [112, 113].

Other contributions to the background for $0\nu\beta\beta$ search come from ^{210}Tl and ^{208}Tl . However, these isotopes are created by α decays of ^{212}Bi and ^{214}Bi , which have smaller branching ratio compared to the β decay (0.021% and 36% respectively [64]). Thus the expected number of these nuclei in the detector at any given time is smaller compared to isotopes higher in the decay chain (especially in the case of ^{210}Tl). Finally, a small portion of $^{234\text{m}}\text{Pa}$ and ^{228}Ac β decays will be smeared to the signal region by energy resolution. All these decays are accompanied by the emission of γ s and, therefore, in addition to purification and coincidence tagging, internal backgrounds can also be reduced by the multi-site background discrimination technique, which was explicitly demonstrated in [102].

6.1.4 Cosmogenic Isotopes

Cosmogenic isotopes (or "cosmogenics" for short) are created via cosmic ray spallation, a reaction of cosmic rays with an atomic nucleus that leads to a creation of a daughter nucleus different from its parent. The resulting isotope can be radioactive and its decay can act as a background in SNO+.

For the $0\nu\beta\beta$ search, these backgrounds are produced by spallation reaction on tellurium [60]. This makes them particularly dangerous as they scale with tellurium loading in the same way as is expected from the signal. Furthermore, decays of some of these isotopes can appear as a single energy peak in the vicinity of the signal region.

The low cosmic ray flux in SNOLAB means that the production of cosmogenics in the laboratory is very small. However, the telluric acid used in SNO+ was produced and transported on the Earth's surface and, therefore, presence of cosmogenic nuclides in the material is expected. As described in section 4.6, the strategy for mitigation of cosmogenic backgrounds is the combination of underground storage of the telluric acid prior to deployment and removal of remaining contaminants through purification. Consequently, less than one event per year is expected in the ROI due to cosmogenic background. However, the aforementioned worrisome properties of these backgrounds require additional caution in case of unexpected contamination, which is not removed by the purification process. Isotopes that decay on the timescale of years or are fed by long lived-parents are of particular concern. A list of potentially problematic isotopes that can be produced by cosmogenic neutron and proton spallation was compiled in [114]. Table 6.2 lists these isotopes with Q-value > 2.3 MeV together with their half-lives. In addition, ^{130}I and its excited states can be created by solar neutrino capture on natural tellurium [115]. These isotopes are relatively short-lived, but can be created inside the detector during data taking as the neutrino flux is not significantly affected by the underground location.

Isotope	Q-value [MeV]	Half-life [days] (parent)
^{22}Na	2.84	950.6
^{26}Al	4.00	2.62×10^8
^{42}K	3.53	0.51 (1.2×10^4)
^{44}Sc	3.65	0.17 (2.16×10^4)
^{46}Sc	2.37	83.79
^{56}Co	4.57	77.2
^{58}Co	2.31	70.9
^{60}Co	2.82	1925.27 (5.48×10^8)
^{68}Ga	2.92	4.7×10^{-2} (271)
^{82}Rb	4.40	8.75×10^{-4} (25.35)
^{84}Rb	2.69	32.8
^{88}Y	3.62	106.63 (83.4)
^{102m}Rh	2.32	1366.77
^{102}Rh	2.46	207.3 (1366.77)
^{106}Rh	3.54	3.47×10^{-4} (371.8)
^{110m}Ag	3.01	249.83
^{110}Ag	2.89	2.85×10^{-4} (249.83)
^{124}Sb	2.90	60.2
^{126m}Sb	3.69	0.01 (8.4×10^7)
^{126}Sb	3.67	12.35 (0.01)
^{130m}I	2.99	6.1×10^{-3}
^{130}I	2.95	0.515

Table 6.2: Potential cosmogenic backgrounds with Q-value > 2.3 MeV. Only isotopes produced by proton and neutron spallation with direct half-life > 20 days are considered or ones that are daughters of longer-lived parents (half-life in brackets in the last column). On top of the isotopes identified in [114], ^{130}I , which is produced by neutrino capture on tellurium, was added to this table. Data taken from [64].

Figure 6.6 shows the reconstructed energy distribution of the cosmogenics listed in Table 6.2 together with the $0\nu\beta\beta$ peak. It is noted that the decays most similar to the signal are the ones of ^{22}Na , ^{44}Sc , ^{60}Co , ^{88}Y , and ^{124}Sb , which have energy peaks close to $Q_{\beta\beta}$. Other listed isotopes can still contribute to the background, but the energy resolution is sufficient to distinguish between the signal and an unexpected contamination by these isotopes.

Decays of cosmogenic isotopes can be identified using the timing classifier as demonstrated in section 5.3.1. It is shown later in this work (section 6.10) that with this pulse shape discrimination as an additional observable in the signal extraction, cosmogenic contamination cannot be mistaken for a $0\nu\beta\beta$ discovery even if the energy and position distributions are alike.

6.1.5 External Backgrounds

The PMTs, external water, hold-down and hold-up ropes, and the AV itself all have lower radiopurity compared to the detector target. The products of radioactive decays happening outside of the scintillator volume can enter the AV, deposit their energy there and imitate $0\nu\beta\beta$ signal. Alternatively, the detector position resolution may cause them to reconstruct inside the scintillator even if all the energy was deposited outside of it.

Due to the different mean path length of the various primary particles, it is predominantly the γ s accompanying radioactive decays that penetrate and reconstruct deep inside the scintillator volume. The main contributions to the external backgrounds are ^{208}Tl and ^{214}Bi decays. The former produces a 2.6 MeV γ that can mimic the $0\nu\beta\beta$ signal, while the latter results in multiple γ s, whose energies can combine and reconstruct near the $Q_{\beta\beta}$.

As the path length of the γ in scintillator is attenuated exponentially, the probability of an external event reconstructing at a certain radial position also falls exponentially towards the detector centre. This is demonstrated in Figure




Figure 6.6: The reconstructed energy distributions of the potentially problematic cosmogenic isotopes from Table 6.2 as simulated by RAT. The investigated $0\nu\beta\beta$ peak at $Q_{\beta\beta} = 2.527$ MeV is shown in black for comparison. Curves are normalised to unity.

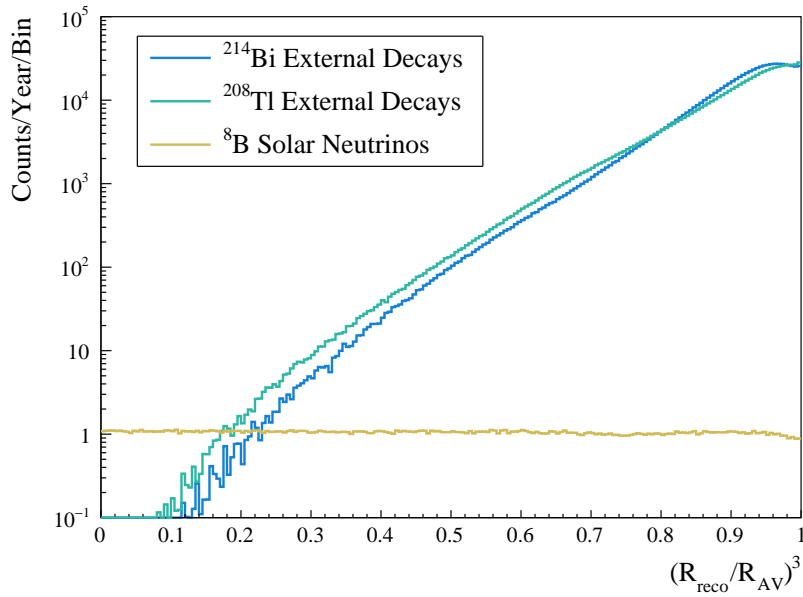


Figure 6.7: Radial distributions of external ^{208}Tl (green) and ^{214}Bi (blue) decays. Horizontal axis gives the volume-corrected reconstructed radius, where R_{reco} and R_{AV} are the event's reconstructed radial position and the radius of the AV, respectively. Only events with reconstructed energies between 1.8 and 3 MeV are shown. Contributions from the hold-down and hold-up ropes, the external water, the AV acrylic and the AV dust are included. Histograms are normalised to expected counts per year. The expected contribution from ^8B neutrino scattering events (orange) is also plotted for comparison.

6.7.

Typically, the main strategy for external background mitigation is fiducialisation, where only events inside certain "fiducial volume" are used in the analysis. However, decreasing this volume also decreases the isotope exposure and, therefore, the fiducial cut needs to be carefully optimised. For counting analyses, this volume is traditionally set to be within the radius, where the external background rate is equal to the rate of ^8B neutrino elastic scattering background. However, for spectral analyses, such as the one presented in this chapter, the dependence of external background rate on radius can be used to constrain this background. For this reason, the radial distance from the detector centre is one of the core observables in this work and the selected fiducial volume is larger (5.5 m).

Importantly, the external backgrounds can also be mitigated using the timing and topological discrimination described in sections 5.3.1 and 5.3.2.

6.1.6 AV Leaching and (α -n) Reactions

After the SNO experiment finished data taking, the cavity and the AV were drained and exposed to air for several years before being repurposed for SNO+. During this time, radon daughters from air were embedded into the AV surface. As shown in Figure 6.4, ^{222}Rn decays to ^{210}Pb , which has a half-life of 22 years and, hence, a significant portion of it will have remained in the material to this day. When the acrylic is put in contact with a liquid, the heavy metals can leach into the liquid. Consequently, there is an additional radioactivity source, feeding the scintillator with ^{210}Pb and its daughters. The specific rate of leaching is highly dependant on the type of materials in contact and the temperature [65]. Various leaching models are investigated and will ultimately be measured in situ. Luckily, all decays following ^{210}Pb in the ^{238}U chain either have Q-values significantly below the ROI or, in the case of α decays, have the visible energies quenched below this level. Therefore, the leaching backgrounds are unlikely to pose as a direct background to $0\nu\beta\beta$.

However, AV leaching is expected to be a major source of α s, particularly from ^{210}Po decay, which can give rise to (α , n) reactions. In these processes, an α interacts with a nucleus to produce a free neutron. In organic scintillator, the most common such reaction is



An (α , n) interaction produces two signals in the SNO+ detector. The prompt signal consists of the scintillation light produced by the α before it captures and by the daughter neutron scattering off of protons in the material. Once the neut-




Figure 6.8: Reconstructed energy distribution of (α, n) interactions before (blue) and after (orange) coincidence cut. Histograms are normalised to expected counts per year taken from [88].

ron is thermalised, it captures on another nucleus, which is left in an excited state. The subsequent de-excitation through γ emission happens on the timescale of $\mathcal{O}(100 \mu\text{s})$ and creates the delayed signal [60]. In SNO+, the neutron captures on hydrogen nuclei in 99% of cases, which de-excites via a 2.2 MeV γ . Consequently, either the prompt or delayed signal can reconstruct in the $0\nu\beta\beta$ ROI. Due to the two signals, however, (α, n) reactions can be rejected using a coincidence tag similar to the one used for rejection of $^{212/214}\text{BiPo}$ out-of-window events. In this work, a selection described in [116] was applied, removing every pair of subsequent events separated by less than 1 ms and with individual reconstructed energies > 1 MeV. Figure 6.8 shows the reconstructed energy spectrum of (α, n) interactions before and after the cut. The number of (α, n) counts per year will depend on ^{210}Po contamination in tellurium and LAB, the leaching rate, duration of the water and unloaded scintillator phases as well as how often is the scintillator recirculated through the purification plant. In this work, the rate calculated in [88] is used.

6.2 Analysis Framework

The purpose of the signal extraction analysis discussed in this chapter is to obtain the most probable value of $0\nu\beta\beta$ half-life given a dataset. In the Bayesian interpretation of statistics, this probability represents the degree of belief in a hypothesis, which in this case is the half-life value or a limit on it. The Bayes' theorem states

$$\underbrace{P(\text{hypothesis}|\text{dataset})}_{\text{posterior probability}} \propto \underbrace{P(\text{dataset}|\text{hypothesis})}_{\text{likelihood}} \cdot \underbrace{P(\text{hypothesis})}_{\text{prior probability}}. \quad (6.5)$$

The posterior is a conditional probability of a hypothesis with certain parameter values given the experimental data. The set of parameters that maximise the posterior is the result of the analysis. In this case, the desired parameter is the number of $0\nu\beta\beta$ counts, but other parameters are also floated in the fit as described below. The likelihood is the probability of observing the data given the hypothesis and the prior contains any existing knowledge about the parameters before the measurement.

In order to fit a hypothesis to the SNO+ data, a Monte Carlo simulation is used to obtain the expected distributions of the individual backgrounds and the signal in the desired observables (energy, position, classifier values etc.). These are the probability distribution functions (PDFs). The parameters optimised in the fit are the numbers of events assigned to each of the backgrounds and the signal, i.e. the normalisations of the PDFs. In this work, ~ 30 parameters are allowed to vary in the fit, which includes the normalisations for all backgrounds with expected rates > 1 count/year after passing the analysis cuts, as well as the potentially problematic cosmogenics (^{60}Co , ^{88}Y , ^{22}Na) and the signal. The full list of fit parameters is included in Appendix F.

The analysis uses the OX0 signal extraction software discussed in [101]. This code is responsible for optimising a test statistic given a dataset, PDFs and a

set of parameters and cut values. Both the dataset and PDFs must have the same dimensionality and axes, meaning the distributions must be in the same observables. The selected test statistic is the extended binned log-likelihood function:

$$\log \mathcal{L}(\vec{\mu}) = - \sum_{i=1}^{N_{par}} \mu_i + \sum_{j=1}^{N_{bin}} n^j \log \left(\sum_{i=1}^{N_{par}} \mu_i f_i^j \right), \quad (6.6)$$

where $\vec{\mu}$ is a vector of N_{par} parameters with μ_i giving the number of events of type i [103]. With both the dataset and PDFs having N_{bin} bins, n^j is the number of observed events in bin j and f_i^j is the content of bin j of the PDF corresponding to event class i .

Prior probabilities (simply priors) are useful for including constraints from other measurements, analyses and calibration. Where such constraints are not available, uniform priors are used, which is the case in this work. The only exception is the normalisation of ${}^8\text{B}$ solar neutrino scattering, where a Gaussian prior calculated in [101] was used, centred on the value coming from the flux measurement by the SNO experiment [10]. The priors are added directly to the $\log \mathcal{L}$ in OX0.

The likelihood has dimensionality equal to the number of parameters (~ 30) and, therefore, a robust optimisation algorithm is needed for the fit. OX0 offers a choice of optimisation routines, including an implementation of Hamiltonian Dynamics Markov Chain Monte Carlo (HMC). As any other Markov Chain Monte Carlo (MCMC), it obtains the posterior distribution by sampling the probability space and accepting or rejecting each step depending on the posterior value at the proposed step relative to the current one. In this case, it is sampling the likelihood space including any non-uniform priors. The acceptance conditions are set so that the samples with higher posterior probability value have higher chance of being accepted. Compared to other versions of MCMC, HMC uses Hamiltonian dynamics to propose the steps in the chain, which results in faster convergence in multi-dimensional probability spaces with correlated para-

meters. A more detailed description of the algorithm can be found in [117].

Finally, it is necessary to extract the posterior probability for $0\nu\beta\beta$ signal from the overall multi-dimensional distribution. This is done by simply integrating the posterior probability over all parameters except for $0\nu\beta\beta$ normalisation. The process is called marginalisation and is also included in the framework. With $0\nu\beta\beta$ counts as random variable and only flat or Gaussian priors, the marginalised posterior probability is expected to follow a Gaussian distribution. However, non-Gaussian tails can arise due to correlation with other parameters.

The creation and choice of the components that are input into the analysis (datasets, PDFs, cuts and observables) are described in the following four sections.

6.3 Monte Carlo Simulation and Fake Datasets

The studies in this chapter use an official SNO+ Monte Carlo production simulated and reconstructed with RAT 6.16.3 with 0.5 % Te-loaded scintillator cocktail including DDA as described in Chapter 4. Additional simulation was run locally on a computer cluster at University of Oxford with the same version of the software to increase the statistics on the following backgrounds: internal ^{208}Tl , ^{234m}Pa , $^{214}\text{BiPo}$ and most of the external backgrounds. All background sources listed in section 6.1, which are expected to contribute to the reconstructed energy spectrum around $Q_{\beta\beta}$, are simulated separately with the exception of $^{212/214}\text{BiPo}$ and (α, n) interactions. In these cases, the entire interaction or part of the decay chain is simulated. The $^{212/214}\text{BiPo}$ and (α, n) coincidence cuts were applied as described in sections 6.1.3 and 6.1.6 as a first step of the analysis.

The SNO+ experiment has not started its tellurium phase at the time of writing of this work. Therefore, the studies presented here use datasets from MC rather than real data. Two types of datasets are used in this chapter: Asimov and standard sample datasets.

Asimov datasets are a convenient way for obtaining the median sensitivity of an experiment [118]. They are representative datasets, which are created with a large event sample and are scaled down to the expected normalisation. As a result, an Asimov dataset does not include fluctuations in its distribution. Instead it is the sum of the average distributions of the individual backgrounds. To scale each of these contributions, the average distributions are normalised to the counts expected in the full detector for the particular background during a set data taking period.

The standard sample datasets, on the other hand, are created simply by drawing only the expected number of events for the given event type (including Poisson fluctuation in this case). As a result, these datasets are individually not representative of the experiment's sensitivity and include bin to bin fluctuations. An ensemble of individual sample datasets can be used to serve the same purpose as an Asimov dataset, however, this technique is computationally more demanding as a lot of MC is required and each dataset must be fitted individually. Nevertheless, a small ensemble of 16 such datasets is used in this work in order to verify the robustness of the analysis.

In order to ensure independence of the datasets and PDFs, the simulation was always split into two parts: one used for obtaining the PDFs and one for creating fake datasets. For studies with Asimov datasets, the MC was split into two independent halves and one of them is used entirely to create one Asimov dataset. For the standard sample datasets, the situation is more complicated due to lower statistics on the external backgrounds. In total, 16 standard sample datasets were created. Four datasets are fully independent with no event repetition. The remaining 12 datasets have external background events drawn randomly from the same pool of events used in the first four datasets. Therefore, these are considered "semi-independent" from each other as well as the first four datasets. The PDFs used in fits with standard sample datasets contain all of the events

$\langle m_{\beta\beta} \rangle$ [meV]	0	25	50	75	100	150	200
$0\nu\beta\beta$ counts in SNO+ [year ⁻¹]	0	6.22	24.87	55.96	99.48	223.83	397.91

Table 6.3: Number of $0\nu\beta\beta$ events in Asimov datasets per year of live-time. The number of ^{130}Te nuclei used in the calculation is $N_{\beta\beta} = 6.25 \times 10^{27}$, corresponding to 0.5 % loading of natural tellurium by weight.

not utilised for creation of the given dataset. It was ensured that at least the equivalent of 10 years of MC is available for each event type for PDFs creation.

All of the datasets assume a live-time of 3 years. The rates assumed for the various background contributions were taken from [88] and are listed in Appendix F. Cosmogenic background rates vary with time spent underground and, therefore, the average rate of the first five years following a year of cool down was used. For the studies presented in this chapter, variable strengths of the $0\nu\beta\beta$ signal were inserted to Asimov datasets corresponding to different effective Majorana masses. Table 6.3 gives the number of $0\nu\beta\beta$ events² in the entire detector volume corresponding to different values of $\langle m_{\beta\beta} \rangle$. The standard sample datasets were created with no signal and $\langle m_{\beta\beta} \rangle = 150$ meV scenarios only. The values of the phase-space factor and nuclear matrix element used for the conversion of $\langle m_{\beta\beta} \rangle$ to $T_{1/2}^{0\nu}$ were $G_{0\nu} = 3.69 \times 10^{-14}$ years⁻¹ and $M_{0\nu} = 4.03$ [119,120]. The results presented later in this chapter in sections 6.8 and 6.9 were converted back to half-life using the same values to provide model-independent sensitivity estimates.

6.4 Observables

An observable is a variable measured for each event. It can either be an information coming directly from the detector electronics or be a result of processing the event's data with a software routine (such as a reconstruction algorithm or

²additional datasets were created for Figure 6.16, but the same calculation was used

a classifier). As the signature for $0\nu\beta\beta$ is a peak in the energy spectrum, reconstructed energy, E , is the primary observable for all searches for this process. In SNO+, the energy resolution of $\sim 3\%$ demonstrated in section 5.2.2 is sufficient to reject most of internal radioactive backgrounds. To deal with external radioactive backgrounds, which have a non-uniform radial distribution, another observable is employed. A reconstructed volume-corrected radius,

$$R_{\text{eff}} = \left(\frac{R_{\text{reco}}}{6000 \text{ mm}} \right)^3, \quad (6.7)$$

is used, where R_{reco} is the distance between the event and the detector centre in units of mm reconstructed by `ScintFitter` and 6000 mm is the radius of the AV. While the signal and internal backgrounds are expected to have a flat distribution in this observable, the external backgrounds are peaked towards high values of R_{eff} .

As demonstrated in Chapter 5, some multi-site backgrounds can be discriminated using the timing and topological classifiers, but cannot be rejected on an event to event basis. In particular, decays of some cosmogenic isotopes might have distributions coincident with $0\nu\beta\beta$ signal in both energy and R_{eff} , but not the timing classifier discriminant. It is shown in this work that adding the output of the multi-site classifiers, introduced in section 5.3, into the list of observables can break the degeneracy between the signal and the cosmogenic backgrounds. The observables are named PSD_{time} and PSD_{topo} for the timing and topological classifiers, respectively. The correlation between individual backgrounds and the signal with and without the pulse shape discrimination is discussed in more detail in section 6.10.

Measured Variable	Accepted Values	Cumulative Signal Acceptance [%]
Fit Valid	True	96
Reconstructed Radius	0 - 5500 mm	74.6
Reconstructed Energy	1.8 - 3.0 MeV	74.5
BiPoCumul Classifier	0 - 0.0003	74.4
BiPoLikelihood214 Classifier	0 - 16	74.4
ITR Classifier	0.29 - 0.37	74.2
PSD _{time}	-10 - 12	74.1
PSD _{topo}	-7 - 5	73.9

Table 6.4: Cuts used in the analysis discussed in this chapter. BiPoCumul and BiPoLikelihood214 are the names of the two in-window $^{212/214}\text{BiPo}$ classifiers discussed in section 6.1.3. Cuts are listed in the order in which they are applied and the last column gives the $0\nu\beta\beta$ signal acceptance after the cut.

6.5 Analysis Cuts

Cuts serve to remove parts of the data that are of no interest for the physics analysis. These can be for example products of hardware issues or events that were not reconstructed properly. Importantly, cuts are also used to reject backgrounds. This section describes the "analysis cuts", i.e. cuts that are applied as part of the analysis.

Table 6.4 gives an overview of the cuts used. First, a valid result from ScintFitter is required, meaning both position and energy fits must have converged successfully. This leads to signal loss of 4%, however, it is noted that almost all of the rejected events originated in the 0.5 m shell nearest to the AV, where the reconstruction is worsened by complicated optics. The reconstruction efficiency inside the inner 5.5 m is $> 99\%$.

The boundaries on the reconstructed energy and radius define the region of interest and the axis limits of these observables. The range of accepted values

for both of these variables is larger than what would be used in a counting analysis. This "extended ROI" is used in order to constrain both internal and external backgrounds using their energy and position distributions as discussed in section 6.1.

The accepted values for the PSD parameters also define the axes for these observables and the values were selected by requiring signal loss to not be larger than $\mathcal{O}(0.1\%)$. The cut values for the in-window $^{212/214}\text{BiPo}$ classifiers were taken from [121].

Finally, the In-Time-Ratio (ITR) classifier is used as a measure of goodness of the position fit. For each event, it gives the ratio of PMT hits falling into a predefined range of the reconstructed time residuals ($-2.5 - 5$ ns). This cut ensures that the time residuals have distribution peaked around 0 ns as expected for a correctly reconstructed event. The cut values were again optimised based on the ITR distribution of $0\nu\beta\beta$ events and the requirement of low signal loss.

All of the analysis cuts were applied to both datasets and PDFs prior to fitting. The resulting signal acceptance of $73.90 \pm 0.04\%$ is obtained from the PDF part of the MC, containing $\mathcal{O}(10^5)$ events). The error on the signal efficiency only accounts for statistical fluctuations.

6.6 Probability Density Functions

The simulation used for generating the PDFs is described in section 6.3. Ideally, one would use the MC to create a full multi-dimensional PDFs with axes corresponding to the selected observables. However, with increasing dimensionality, the number of bins in the PDFs geometrically increases. Consequently, the demand on the amount of MC needed to ensure sufficient statistics in each bin also grows rapidly. This section describes the construction of the PDFs and the selected strategy for dealing with higher dimensionality when including the PSD parameters as additional observables.

The energy distribution contains the principal information for distinguishing the signal peak from the background and, hence, fine binning in energy is needed. On the other hand, no additional knowledge can be gained from the data by including bins much smaller than the energy resolution. Therefore, the width of the energy bins was set to be of the same order as the resolution at $Q_{\beta\beta}$ and the extended ROI was split into 48 bins of 25 keV width. For the remaining observables, the bin count was intentionally kept low, while still allowing the underlying distributions in these observables to be evident. The volume corrected radius was split into 6 bins and the PSD parameters into 5 bins each.

With this binning, the full four-dimensional PDFs of the form

$$P(E, R_{\text{eff}}, \text{PSD}_{\text{time}}, \text{PSD}_{\text{topo}}) \quad (6.8)$$

contain 72000 bins. Unfortunately, the MC was insufficient to provide enough statistics to obtain acceptable PDFs of this form. When tested, some bins were empty across all of the distributions and, hence, would yield infinities in the log-likelihood calculations. This was not remedied even by applying smoothing to some of the lower-statistics PDFs. An alternative approach would be to substitute some parts of the distributions by analytical predictions, however, this method was not explored here. Instead, "PDF factorisation" was employed. If all of the observables were independent from each other, the PDF in equation 6.8 would be equivalent to the product $P(E) \times P(R_{\text{eff}}) \times P(\text{PSD}_{\text{time}}) \times P(\text{PSD}_{\text{topo}})$. However, this is not the case as the observables are correlated with each other and simply factorising all of them could lead to biases in the fit.

It was described in section 5.2.2 that the energy reconstruction algorithm takes the reconstructed position as an input and that the energy bias and resolution depend on position. Furthermore, both topological and timing discriminants rely on the multi-site features and, therefore, they are expected to be correlated with each other. This has been shown explicitly in Figure 5.13. In fact, there

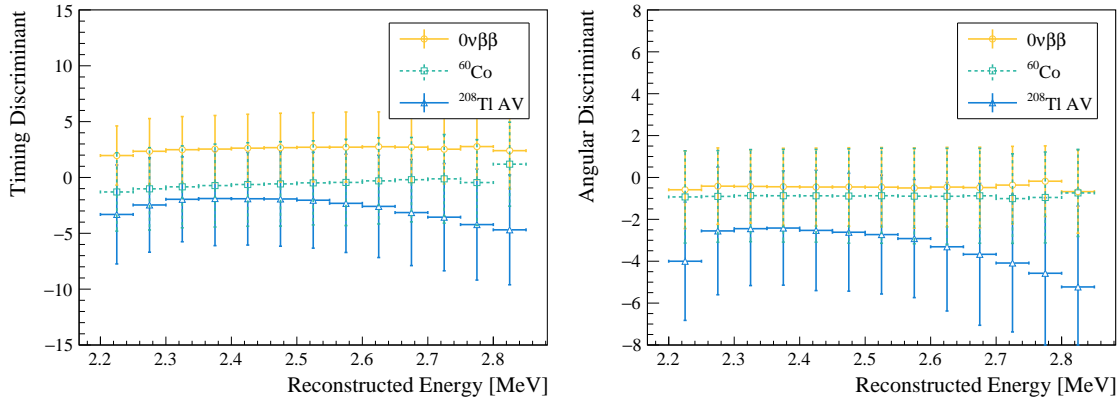


Figure 6.9: Dependence of timing (left) and topological (right) PSD discriminants on reconstructed energy for $0\nu\beta\beta$ (orange circles), ^{60}Co (green squares) and external $^{208}\text{Tl AV}$ events (blue triangles). Points and error bars show the mean and standard deviation, respectively, for the discriminant distribution in each bin.

is correlation between the two classifier parameters for the external multi-site backgrounds but not the signal, which can be helpful in breaking the degeneracy between the background and signal normalisations. This information would be lost if the two PSD observables were inappropriately factorised.

Figures 6.9 and 6.10 show the dependence of the PSD variables on E and R_{eff} , respectively. Again, there is a weak dependence of these multi-site discrimination parameters on the other two observables. This is mainly because the time residuals are calculated using the reconstructed position and the vertex resolution depends on both energy and radius. Furthermore, the classifiers were tuned on a sample of events with a certain energy and position as described in Appendix E.

For this work, the dependence of the PSD observables on energy was neglected. This was justified by the reasonably small variation in the discriminant values across the ROI energy range. The dependence on R_{eff} was retained in the form of conditional PDFs. The resulting factorised four-dimensional PDF considered in this chapter has the form:

$$P(E, R_{\text{eff}}) \times P([\text{PSD}_{\text{time}}, \text{PSD}_{\text{topo}}] | R_{\text{eff}}), \quad (6.9)$$

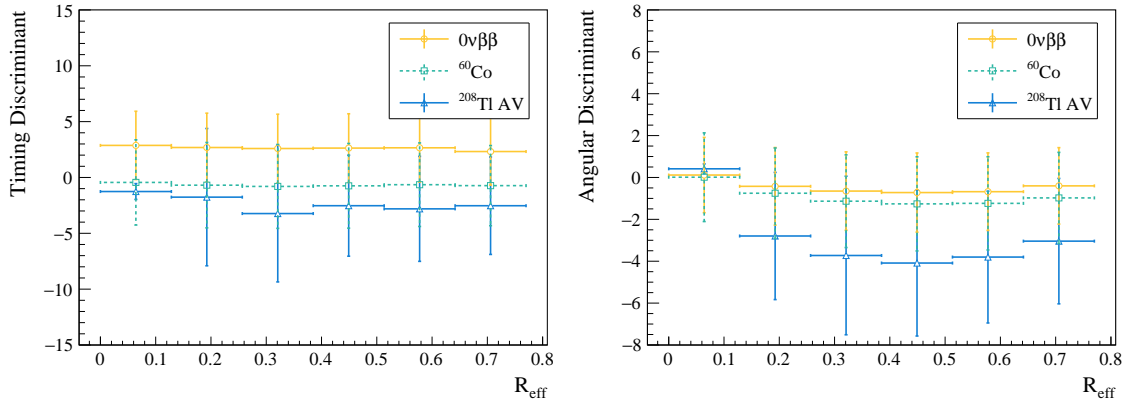


Figure 6.10: Dependence of timing (left) and topological (right) PSD discriminants on reconstructed radial position for $0\nu\beta\beta$ (orange circles), ^{60}Co (green squares) and external ^{208}Tl events (blue triangles). Points and error bars show the mean and standard deviation, respectively, for the discriminant distribution in each bin.

where the both factors are two-dimensional distributions normalised to unity. The latter is a 2D PDF in the classifier discriminants for the relevant radial bin. The validity of this approximation was tested by the fit bias and pull described in section 6.7.

As the PSD observables help with discriminating the signal from backgrounds, their use is only needed for energies around the signal peak. Therefore, the option of applying the factorisation only in a smaller energy region around the $Q_{\beta\beta}$ was also explored. In this case, the PSD observables were simply set to a normalisation constant outside of this region. However, this treatment yielded sharp discontinuities in the PDFs at the place where the boundary of the smaller energy region was set and prevented the fit from converging.

For comparison, two-dimensional fits without any PSD parameters were also performed, for which PDFs of the form

$$P(E, R_{\text{eff}}) \quad (6.10)$$

were used. Furthermore, the topological classifier only facilitates discrimination

between signal and external backgrounds. Its use is expected to improve the experiment's sensitivity by effectively increasing the fiducial volume. However, it cannot help break degeneracies between the signal and internal backgrounds. Therefore, the analysis was also run with full three-dimensional PDFs:

$$P(E, R_{\text{eff}}, \text{PSD}_{\text{time}}). \quad (6.11)$$

Removing the PSD_{topo} observable improved the statistics in each bin and granted successful convergence of the fit. Finally, the factorised version of the 3D PDF was also considered for completeness:

$$P(E, R_{\text{eff}}) \times P(\text{PSD}_{\text{time}} | R_{\text{eff}}). \quad (6.12)$$

It is noted that in fits with both of these 3D PDFs, the hard cut of the topological classifier parameter was still applied for external background rejection.

6.7 Verification of the Signal Extraction

Before drawing conclusions on sensitivity of SNO+ using this analysis, it is necessary to ensure the validity of the results obtained by the fit. To confirm that the analysis successfully recovers the number of $0\nu\beta\beta$ counts in the dataset within the uncertainty, the ensemble of 16 standard sample datasets with signal was fitted using the signal extraction with the 4 different PDF types. For each dataset, the extracted number of $0\nu\beta\beta$ counts, $C_{0\nu}^{\text{fit}}$, and the uncertainty on this parameter, $\sigma_{0\nu}^{\text{fit}}$, are obtained as the mean and the width of a Gaussian fit to the posterior probability. These numbers can then be compared to the number of $0\nu\beta\beta$ events inserted into the dataset multiplied by the signal efficiency, $C_{0\nu}^{\text{expected}}$.

A common method for quantifying the goodness of a likelihood fit is to obtain

the bias and the pull, defined as:

$$\text{Bias} = \frac{C_{0\nu}^{\text{fit}} - C_{0\nu}^{\text{expected}}}{C_{0\nu}^{\text{expected}}} \quad \text{Pull} = \frac{C_{0\nu}^{\text{fit}} - C_{0\nu}^{\text{expected}}}{\sigma_{0\nu}^{\text{fit}}}. \quad (6.13)$$

In an unbiased fit with uncertainties properly described by $\sigma_{0\nu}^{\text{fit}}$, the pulls from many sample datasets are expected to follow a normal distribution centred at 0 with standard deviation (SD) of 1.

The biases and pulls were calculated for the individual sample datasets and the resulting means and standard deviations on these quantities are recorded in Tables 6.5 and 6.6. It is seen from the large mean bias that the 2D fit underestimates the number of $0\nu\beta\beta$ counts in the datasets. Sections 6.9 and 6.10 will demonstrate that this is due to correlation between the signal and the cosmogenic backgrounds, which results in $0\nu\beta\beta$ counts being assigned to this background. Including the PSD parameters significantly reduces the deviation of the bias from 0. The standard deviation of the pull distribution is consistent with unity for fits including PSD parameters. Furthermore, introduction of bias by the factorisation of the PDFs is not observed within the statistics used.

PDF type	Mean Bias	Bias SD
$P(E, R_{\text{eff}})$	-0.35 ± 0.04	0.16 ± 0.03
$P(E, R_{\text{eff}}, \text{PSD}_{\text{time}})$	-0.08 ± 0.03	0.12 ± 0.02
$P(E, R_{\text{eff}}) \times P(\text{PSD}_{\text{time}} R_{\text{eff}})$	-0.05 ± 0.03	0.11 ± 0.02
$P(E, R_{\text{eff}}) \times P([\text{PSD}_{\text{time}}, \text{PSD}_{\text{topo}}] R_{\text{eff}})$	-0.07 ± 0.07	0.29 ± 0.05

Table 6.5: Mean bias and its standard deviation (SD) from an ensemble of standard sample datasets for different versions of the fit.

A graphical representation of the fit results for the dataset ensemble is also shown in Figure 6.11. In these plots, each point represents the number of $0\nu\beta\beta$ counts obtained as the mean of a Gaussian fit to the posterior probability. The errors give the fitted 1σ width of the distribution. It is expected that $\sim 68\%$ of

PDF type	Mean Pull	Pull SD
$P(E, R_{\text{eff}})$	-1.5 ± 0.2	0.7 ± 0.1
$P(E, R_{\text{eff}}, \text{PSD}_{\text{time}})$	-0.7 ± 0.2	0.9 ± 0.2
$P(E, R_{\text{eff}}) \times P(\text{PSD}_{\text{time}} R_{\text{eff}})$	-0.4 ± 0.2	0.9 ± 0.2
$P(E, R_{\text{eff}}) \times P([\text{PSD}_{\text{time}}, \text{PSD}_{\text{topo}}] R_{\text{eff}})$	-0.3 ± 0.2	0.8 ± 0.2

Table 6.6: Mean pull and its standard deviation (SD) from an ensemble of standard sample datasets for different versions of the fit.

the these error bars should overlap with the true value, indicated by the dashed line in each plot. While the fitted number of $0\nu\beta\beta$ counts is again lower than expected in the 2D version of the analysis, fits including the pulse shape dimensions comply with the prediction well. It is concluded that within the available statistics, the bias in the fit with PSD observables is minimal, especially when compared to the 2D signal extraction.

6.8 Limit on $0\nu\beta\beta$ Half-life

In order to determine what limit on $T_{1/2}^{0\nu\beta\beta}$ will SNO+ be able to set after 3 years of data taking, the signal extraction was run on an Asimov dataset with no signal counts inserted into it. The resulting marginalised posterior probability for $0\nu\beta\beta$ normalisation, $C_{0\nu}$, is shown in Figure 6.12. The four different curves correspond to running the versions of the signal extraction, each corresponding to one of the PDFs described in section 6.6. As expected, the probability is peaked towards $C_{0\nu} = 0$ in all cases. However, the distribution is the widest for the 2D fit and narrowest for 4D. Consequently, if there are no signal events in the dataset, the 4D fit is able to reject a higher values of the signal normalisation and will give a stronger limit on $0\nu\beta\beta$ half-life at a given credible interval (C.I.).

A 90 % C.I. is quoted here. To obtain a value for the $T_{1/2}^{0\nu}$ limit from the posterior probability, the distribution is integrated from zero to first obtain the

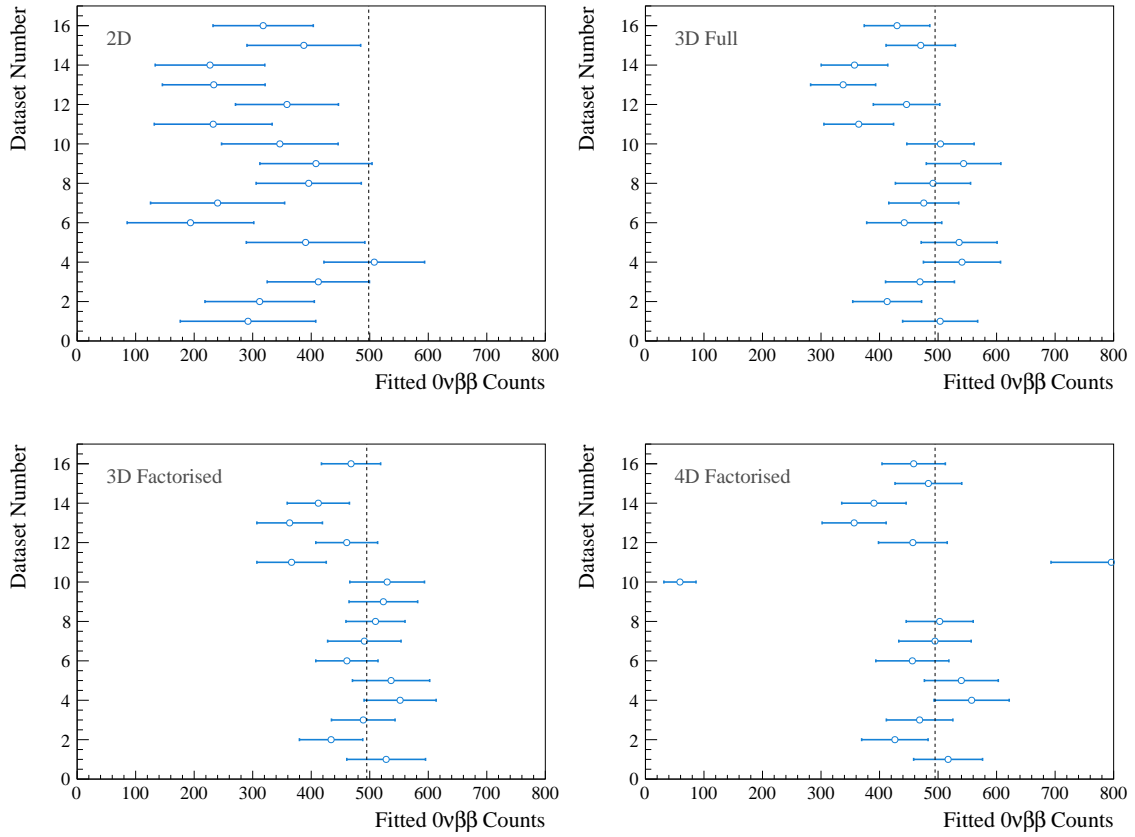


Figure 6.11: Results from running the signal extraction on an ensemble of sample datasets. Individual plots show results from signal extraction ran with 2D PDFs from equation 6.10 (top left), full 3D PDFs from equation 6.11 (top right), factorised 3D PDFs from equation 6.12 (bottom left) and factorised 4D PDFs from equation 6.9 (bottom right). Points and error bars obtained as the mean and σ of a Gaussian fit to the posterior probability from each dataset.

value of $0\nu\beta\beta$ normalisation that can be ruled out at the required confidence, $C_{0\nu}^{90\%}$. The point where the cumulative probability reaches 90 % gives this value, which is then converted to the half-life limit using

$$T_{1/2}^{0\nu}(90\% \text{C.I.}) = \frac{\ln(2)}{C_{0\nu}^{90\%}} \epsilon t N_{\beta\beta}, \quad (6.14)$$

where ϵ is the signal efficiency, t is the live-time and $N_{\beta\beta}$ is the number of ^{130}Te nuclei inside the AV. The resulting half-life limits are reported in Table 6.7.




Figure 6.12: Marginalised posterior probability of $0\nu\beta\beta$ counts from fit to Asimov dataset with no signal. The horizontal axis gives the $0\nu\beta\beta$ normalisation parameter value, which indicates the number of events assigned to the signal. Curves correspond to signal extraction using PDFs in equations 6.10 (orange), 6.11 (green), 6.12 (turquoise) and 6.9 (dark blue). Vertical line indicates the 90% C.I. for the 4D signal extraction.

PDF	$T_{1/2}^{0\nu}$ at 90% C.I. [yrs]
$P(E, R_{\text{eff}})$	$> 1.69 \times 10^{26}$
$P(E, R_{\text{eff}}, \text{PSD}_{\text{time}})$	$> 1.75 \times 10^{26}$
$P(E, R_{\text{eff}}) \times P(\text{PSD}_{\text{time}} R_{\text{eff}})$	$> 1.70 \times 10^{26}$
$P(E, R_{\text{eff}}) \times P([\text{PSD}_{\text{time}}, \text{PSD}_{\text{topo}}] R_{\text{eff}})$	$> 1.80 \times 10^{26}$

Table 6.7: Limits on $T_{1/2}^{0\nu\beta\beta}$ at 90% C.I. for various versions of the fit.

6.9 Discovery Potential

Ideally, SNO+ would like to observe a $0\nu\beta\beta$ signal rather than just set a limit on its half-life. In order to obtain the potential of the experiment to do so, the analysis was run on Asimov datasets with MC signal counts inserted into them. The number of $0\nu\beta\beta$ events in each dataset was summarised in Table 6.3.

Again, the signal extraction with all four versions of PDFs was run over the datasets and the posterior probabilities from the fits are shown in Figure 6.13. Compared to the 2D fits, the posterior probability is shifted towards higher values of signal counts when PSD observables are included. This is again demonstrating the decreased bias in the $0\nu\beta\beta$ count estimation. Importantly, for some values of $\langle m_{\beta\beta} \rangle$ (i.e. 100 and 150 meV) that are consistent with 0 counts in the 2D fit, the posterior probability lies almost entirely above this value. As a result, including the PSD observables make it possible to reject the no signal hypothesis in these cases.

Figure 6.14 shows an example of the fit result. Specifically, the energy projections of the PDFs scaled to the fitted normalisation values are shown for the case of 150 meV signal and analysis with 4D factorised PDFs. Each sub-plot corresponds to a different radial bin. The fake data is also shown. Individual backgrounds listed in Appendix F were grouped into the relevant categories: external backgrounds, internal ^{232}U and ^{238}Th chains, ^8B neutrino scattering, cosmogenic backgrounds, $2\nu\beta\beta$ and (α, n) interactions. It is demonstrated that the dataset is well fitted and that the signal peak is the most visible in the inner most radial bin as expected.

The energy projection in this most central radial bin is also compared with the one resulting from fitting without using PSD in Figure 6.15. It is shown that in the 2D case, the signal counts are incorrectly assigned to the cosmogenic background. The inclusion of the timing classifier discriminant provides the necessary information to distinguish between the signal and the cosmogenics,

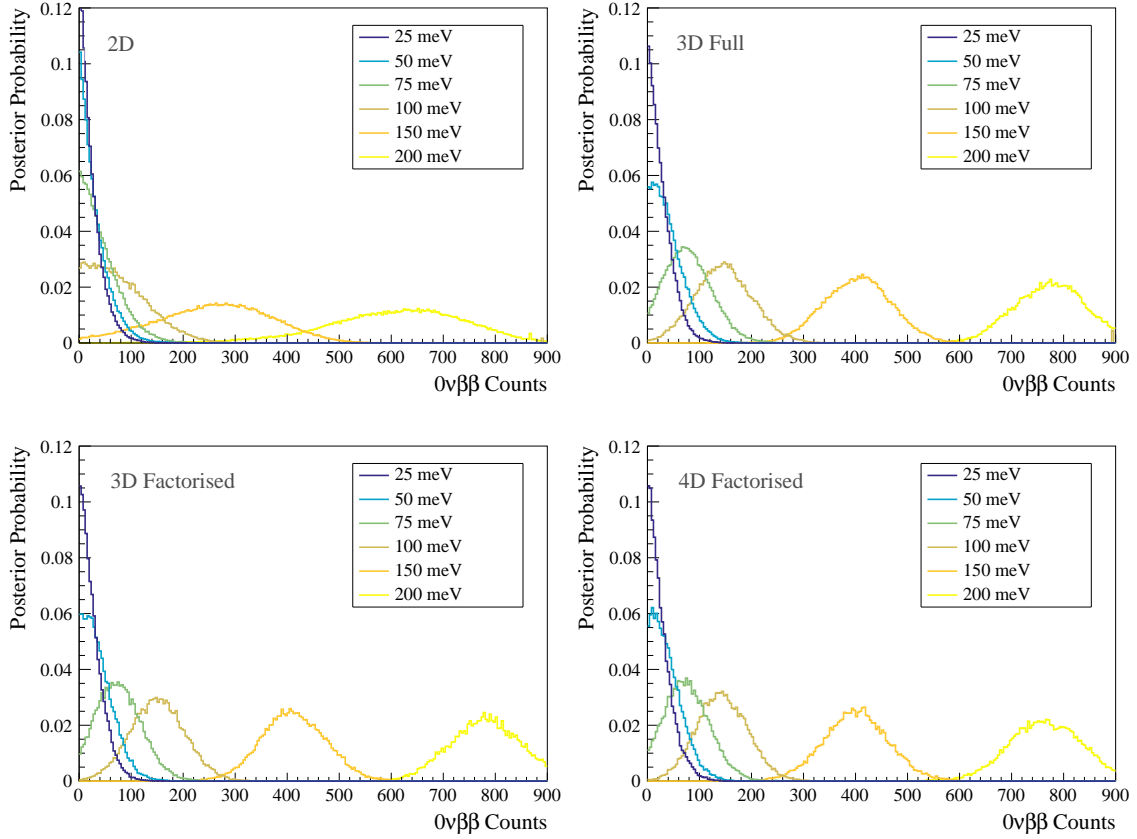


Figure 6.13: Marginalised posterior probability of $0\nu\beta\beta$ counts from fit to Asimov datasets with signal. Individual plots show results from signal extraction run with 2D PDFs from equation 6.10 (top left), full 3D PDFs from equation 6.11 (top right), factorised 3D PDFs from equation 6.12 (bottom left) and factorised 4D PDFs from equation 6.9 (bottom right). The horizontal axis gives the $0\nu\beta\beta$ normalisation parameter value, which indicates the number of events assigned to the signal.

which then results in a better interpretation of the signal counts.

Similarly to the half-life sensitivity, the discovery potential is obtained from the posterior probability. In this case, the question is whether the sensitivity of the experiment is high enough to observe a signal of a given size at a specified confidence. The best fit value is set to be the most probable value of the posterior probability obtained as a mean of a Gaussian fit. The credible interval is then calculated by cumulating probability bins around the best value in order of decreasing probability. Here, the credible interval that does not include $C_{0\nu} = 0$ was calculated, which gives the confidence at which the experiment can make an

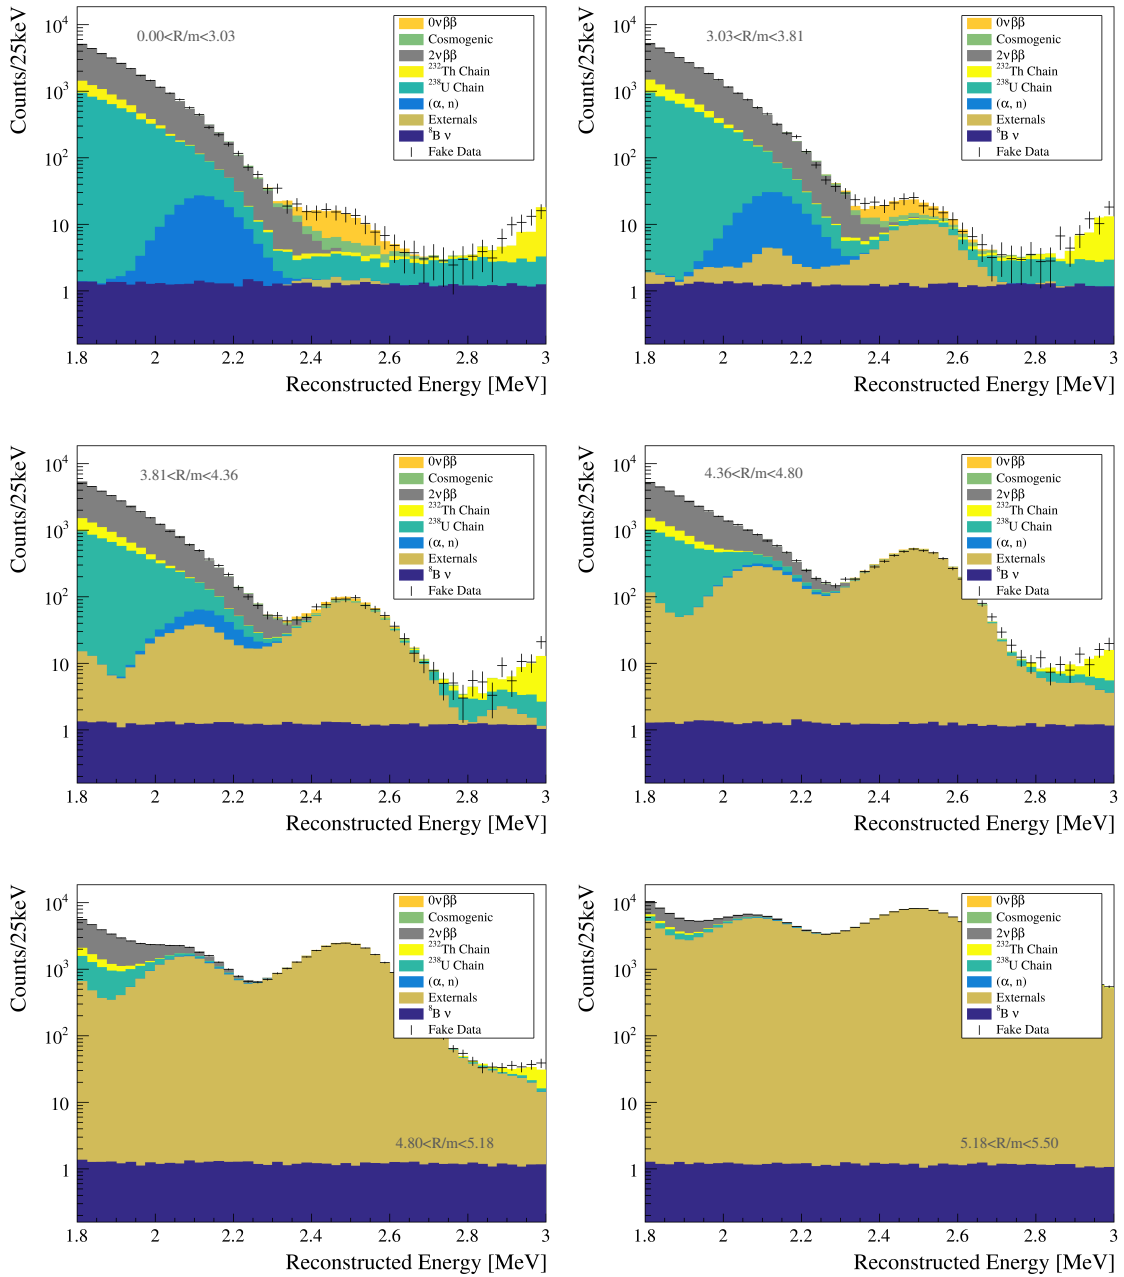


Figure 6.14: An example result from a 4D fit with PDFs from equation 6.9 to an Asimov dataset with $\langle m_{\beta\beta} \rangle = 150$ meV (black points). Stacked distributions of different background groups were scaled to the fitted values. Each plot corresponds to an energy projection in a different radial bin, with top left corresponding to the most central bin and bottom right to the bin closest to the AV.

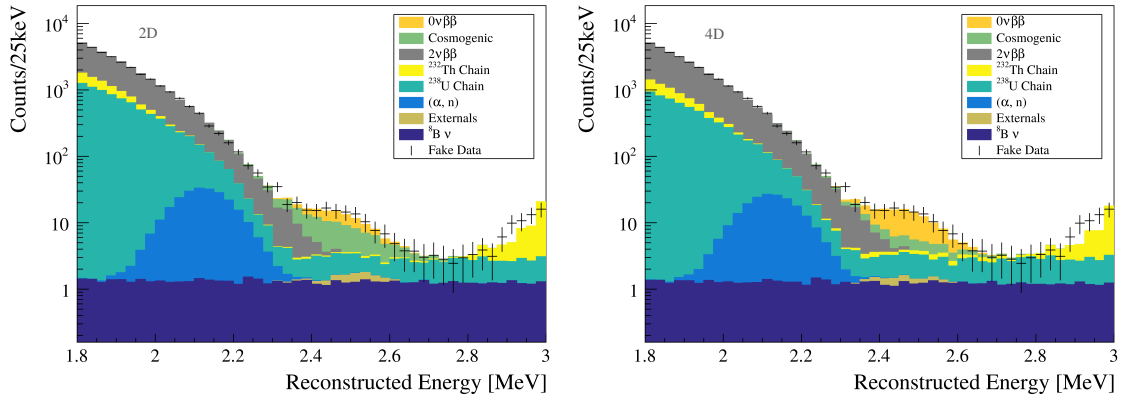


Figure 6.15: Comparison of fit results in the inner most radial bin with 2D (left) and 4D (right) signal extraction. Same dataset with with $\langle m_{\beta\beta} \rangle = 150 \text{ meV}$ (black points) was fitted in both cases. Stacked distributions of different background groups were scaled to the fitted values.




Figure 6.16: Discovery potential for $0\nu\beta\beta$ as a function of $\langle m_{\beta\beta} \rangle$. Dashed and dotted lines indicate 90 % and 99 % C.I., respectively. Individual curves correspond to signal extraction using PDFs in equations 6.10 (orange), 6.11 (green), 6.12 (turquoise) and 6.9 (dark blue).

observation of $0\nu\beta\beta$. This credible intervals are plotted in Figure 6.16 for various signal strengths and signal extraction versions.

The values of nuclear matrix element and phase-space factor used to calculate the $0\nu\beta\beta$ counts in the datasets in Figure 6.16 were stated in section 6.3. To

convert the results from the figure to the effective Majorana mass for different models, the range of matrix element values from [120] was used. It is concluded, that with timing information as an additional observable in the likelihood analysis, SNO+ is expected to have a 99% C.I. discovery potential for $0\nu\beta\beta$ with $\langle m_{\beta\beta} \rangle = 75 - 181$ meV after only 3 years of data taking. Furthermore, SNO+ tellurium phase is planned to run for 5 years. As the sensitivity scales with square root of the live-time, the expected discovery potential is at the level of $\langle m_{\beta\beta} \rangle$ allowed by the current measurements [30,21,26].

Most importantly, the inclusion of PSD observables presents a significant sensitivity gain of 40 – 50 % compared to signal extraction without PSD. The improvement is mainly due to breaking the degeneracy of signal and cosmogenics background PDFs, which is discussed further in the following section.

6.10 Correlation with Cosmogenic Backgrounds

It is argued that the improved sensitivity to $0\nu\beta\beta$ with PSD comes primarily from breaking the correlation between the signal and the cosmogenic backgrounds. This correlation comes from the similarity of these event types in energy and radial observables. It is explicitly demonstrated in Figure 6.17, where two-dimensional posterior probabilities are shown for fits with and without including the timing classifier observable. Three cosmogenic backgrounds are examined in this figure: ^{22}Na , ^{60}Co and ^{88}Y . These isotopes have the most similar energy distributions to $0\nu\beta\beta$ as discussed in section 6.1.4. ^{130}I and ^{130m}I normalisations were also floated in the fit, but the reason for their inclusion as fit parameters was their higher expected rate rather than their ability to mimic the signal. Therefore, the correlation of the posterior probabilities of these two cosmogenics with signal is not shown, but was confirmed to be negligible.

From the three potentially problematic isotopes examined in Figure 6.17, only ^{60}Co was found to have a significant correlation with $0\nu\beta\beta$ in the 2D analysis.

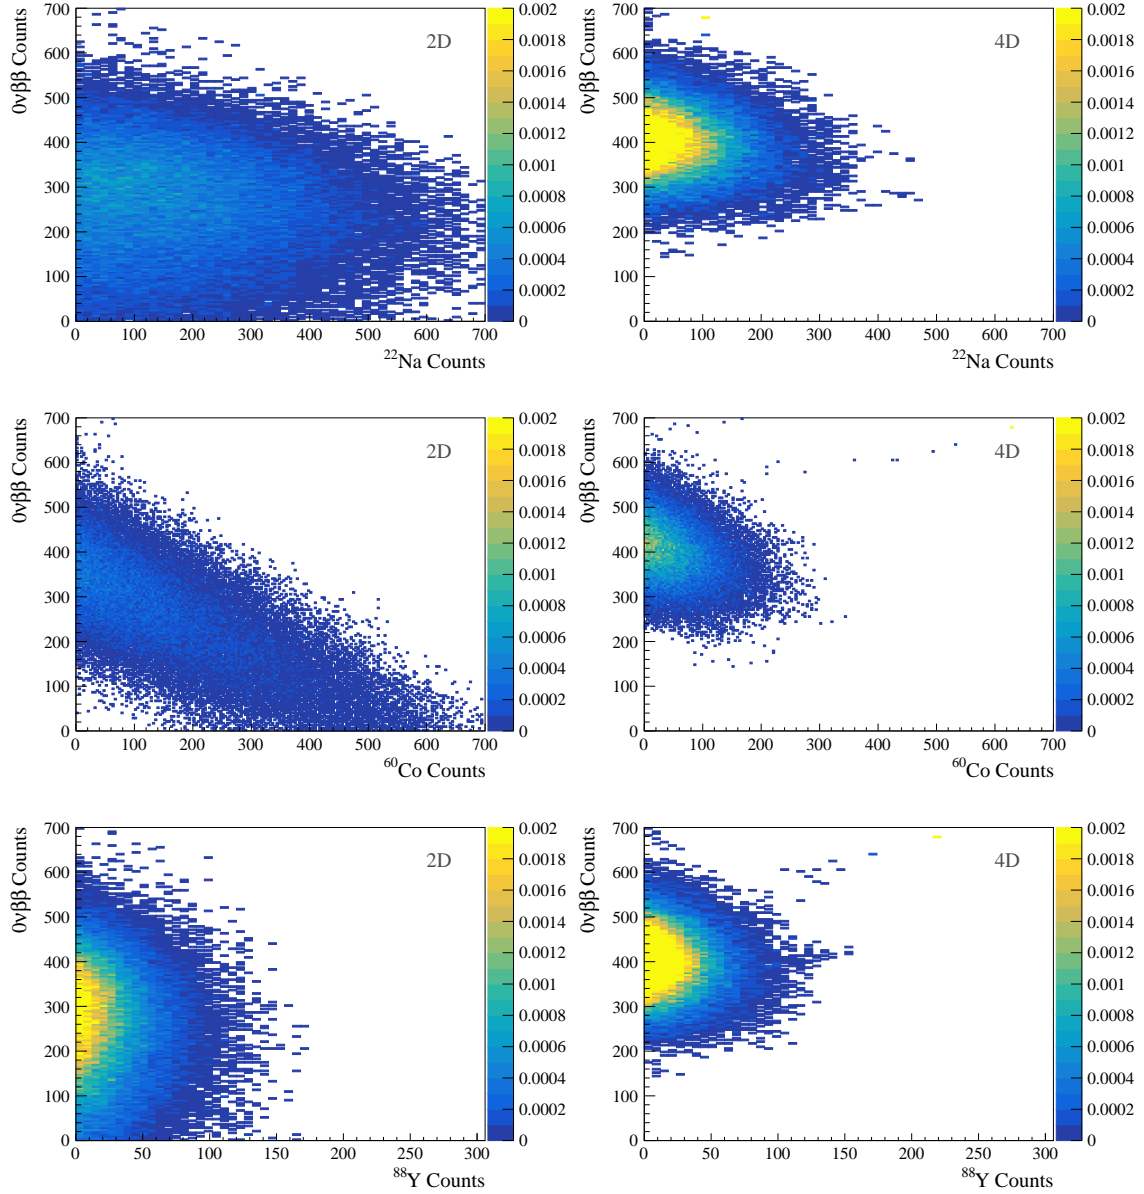


Figure 6.17: Correlation of $0\nu\beta\beta$ posterior probability with ^{22}Na (top), ^{60}Co (middle) and ^{88}Y (bottom). Plots on the left are from analysis without PSD (PDFs from equation 6.10) and plots on the right are from 4D analysis with PSD (PDFs from equation 6.9). The same Asimov dataset with signal corresponding to $\langle m_{\beta\beta} \rangle = 150$ meV was fitted in each case.

Scaling Applied to Expected ^{60}Co Rate	^{60}Co counts	Equivalent $\langle m_{\beta\beta} \rangle$ [meV]
1	3.46×10^{-4}	-
2	6.92×10^{-4}	-
10	3.46×10^{-3}	-
10^2	3.46×10^{-2}	-
10^3	3.46×10^{-1}	-
10^4	3.46×10^1	-
2.87×10^4	9.95×10^1	100
6.47×10^4	2.23×10^2	150
10^5	3.46×10^2	-
1.14×10^5	3.98×10^2	200

Table 6.8: Datasets for studies with increased ^{60}Co background, each with the rate of this background scaled by a different factor given in the first column. The number of counts inserted into a 3 year Asimov dataset with no $0\nu\beta\beta$ signal is indicated in the middle column. For three of these datasets, the scale factor was selected such that the number of ^{60}Co counts is equivalent to number of expected signal counts for a specific $\langle m_{\beta\beta} \rangle$; this is indicated in the last column.

Importantly, this degeneracy is significantly reduced by adding the additional observables into the fit, which is demonstrated on the middle right plot.

To scrutinise the SNO+ sensitivity with regards to the ^{60}Co background further, a study was carried out, where unexpected contamination with this isotope was assumed. In this case, Asimov datasets with live-time of 3 years, no $0\nu\beta\beta$ events and additional ^{60}Co counts were created as summarised in Table 6.8. As the level of this isotope in the detector is predicted to be very low following the purification procedure, an increase of $\mathcal{O}(10^4)$ in the expected rate of this background is needed to create the same number of events that are expected from $0\nu\beta\beta$ with $\langle m_{\beta\beta} \rangle$ in the range tested in this work.

Two versions of the signal extraction were performed on these datasets: a




Figure 6.18: $0\nu\beta\beta$ half-life sensitivity with increased ^{60}Co background with (yellow) and without (purple) PSD. Horizontal axis gives the multiplication factor applied to the expected ^{60}Co rate from Appendix F. Error bars give standard deviation from Monte Carlo sampling (length of MC chain decreased for this study compared to section 6.8).

2D fit without any PSD and a 4D fit with two PSD observables (PDFs in equations 6.10 and 6.8, respectively). The $0\nu\beta\beta$ half-life limit with this additional contamination as a function of the factor used to scale the expected ^{60}Co counts is shown in Figure 6.18. While the achievable half-life expectedly decreases with increasing background, the impact is less damaging when PSD is used in the analysis.

Another crucial test is whether ^{60}Co contamination can be mistaken for $0\nu\beta\beta$ observation. In order to examine this possibility, the posterior probability of $0\nu\beta\beta$ normalisation was examined for the datasets with increased cobalt counts. As no signal was inserted into these datasets, the posterior is expected to be consistent with zero. This is indeed observed in Figure 6.19 for both 2D and 4D analyses. The latter results in a narrower distribution around zero, which is again consistent with the better half-life limit from Figure 6.18. In conclusion, it has been demonstrated, that should the analysis yield a non-zero signal normal-

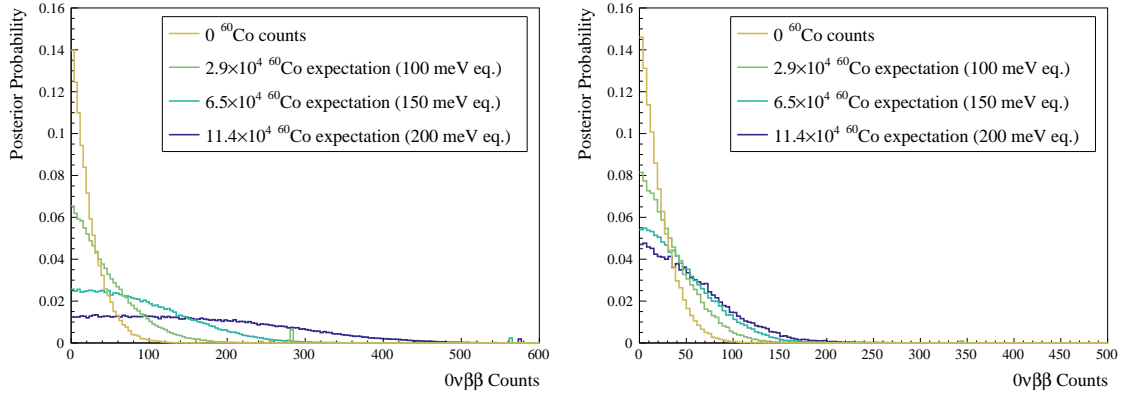


Figure 6.19: Marginalised posterior probability of $0\nu\beta\beta$ counts from fit to Asimov datasets with ^{60}Co contamination and no signal. Results from signal extraction with 2D PDFs from equation 6.10 shown on the left and factorised 4D PDFs from equation 6.9 on the right. The horizontal axis gives the $0\nu\beta\beta$ normalisation parameter value, which indicates the number of events assigned to the signal.

isation at a given confidence, it will be a genuine $0\nu\beta\beta$ observation rather than an unexpected cosmogenic background contamination.

6.11 Summary

In this chapter, a signal extraction for $0\nu\beta\beta$ in SNO+ was presented. Specifically, the discussion regarded the development of a likelihood analysis for this search using Markov Chain Monte Carlo including ~ 30 background event types as parameters and fitting over 4 observables: energy, distance from the centre of the detector, topological and timing pulse shape discrimination parameters. The latter two differentiate between interactions with point-like and multi-site energy deposition and are obtained using a classification described in Chapter 5.

The projected half-life limit and discovery potential for $0\nu\beta\beta$ using this analysis were estimated. In both cases, there was a significant improvement in the expected sensitivity compared to an analysis not including the PSD parameters. Particular gain is observed should there be an unexpected contamination by cosmogenic backgrounds. Ultimately, it was demonstrated that, by including the PSD parameters in the analysis, SNO+ does not only have the potential to set a

world leading limit, but also to make a discovery depending on the true value of effective Majorana mass.

Chapter 7

Conclusions and Outlook

This thesis focused on improving the sensitivity to $0\nu\beta\beta$ in ^{130}Te in SNO+, a liquid scintillator detector. Specifically, the sensitivity was enhanced by two means: modification of the tellurium-loaded scintillator cocktail and the development of a four-dimensional signal extraction including PSD parameters.

The increase of light yield through the addition of DDA was demonstrated in Chapter 4. Moreover, it was shown that DDA can prevent deterioration of the Te-loaded scintillator due to water exposure to the extent that can be tested on a time scale of 3 years. This is particularly important as the scintillator is directly replacing water inside the acrylic vessel and the two liquids will remain in direct contact for a prolonged period of time. The fact that no humidity-induced deterioration of Te-loaded scintillator samples with DDA (at concentrations of $\text{DDA:Te} \geq 0.25 \text{ MR}$) was ever observed significantly enhances the confidence in the tellurium loading technique and the prospect of $0\nu\beta\beta$ search in SNO+. The collaboration now plans to add DDA to the scintillator mixture together with the tellurium loading. Furthermore, the tests presented in this thesis led to further studies with this compound. Preliminary results show that by including DDA during the TeBD synthesis, rather than adding it afterwards, the light yield can be increased even further [90]. This is a promising avenue especially for future

upgrades of the experiment, which would allow for higher tellurium loading.

The algorithms for event reconstruction and classification in SNO+ were discussed in Chapter 5, with a particular focus on the implementation of timing and topological classifiers into the SNO+ software. These classifiers are now being used by the collaboration in multiple physics analyses as a tool for external background rejection.

The outputs of the classifiers were then incorporated as additional dimensions into a likelihood analysis for $0\nu\beta\beta$ search, described in Chapter 6. Compared to an analysis without these dimensions, significant enhancement in the expected sensitivity to the $0\nu\beta\beta$ process was demonstrated. Additionally, it was shown that none of the considered backgrounds can successfully mimic $0\nu\beta\beta$ when the likelihood analysis with PSD is employed. Therefore, should a signal be observed in SNO+ using this analysis, it will be a genuine $0\nu\beta\beta$ observation rather than contamination by radioactive backgrounds. Further improvements to the analysis, which were not explored in this work, could include removing the correlation between the observables. This could be achieved by incorporating energy and position into the calculation of classifier discriminant as well as the tuning of classifier weights. Alternatively, the correlation of the observables could be estimated from an independent MC dataset and a correction could be applied to the observables to create new independent dimensions of the signal extraction. Such treatments would permit a more extensive factorisation of PDFs, which could in turn allow for even more observables in the signal extraction (e.g. the time since the start of data taking, which could help further constrain the cosmogenic backgrounds that are expected to decay within the data taking period). An alternative approach to deal with limited statistics would be to replace some of the PDFs with analytical functions.

With both of the above improvements, the expected lower limit on $0\nu\beta\beta$ was estimated to reach 1.8×10^{26} years at 90 % confidence after 3 years if no signal

is observed, which is higher than any currently set limit. Furthermore, SNO+ is expected to have 99 % C.I. discovery potential for an effective Majorana mass of 75 – 181 meV after the same data taking period. It is concluded that SNO+ will be able to probe $0\nu\beta\beta$ at effective Majorana masses that are 40 – 50 % lower than what is expected without the improvements presented in this thesis.

Appendix A

List of Acronyms

$0\nu\beta\beta$ Neutrinoless double beta decay

$2\nu\beta\beta$ Two neutrino double beta decay

ADC Analog-to-digital converter

AV Acrylic vessel

BD 1,2-Butanediol

bisMSB 1,4-Bis(2-methylstyryl)benzene

BSM Beyond the Standard Model

C.I. Credible interval

DDA N,N-Dimethyldodecylamine

ELLIE Embedded LED/Laser Light Injection Entity

FEC Front End Card

GT Global Trigger

HDR Hold-down ropes

HMC Hamiltonian Dynamics Markov Chain Monte Carlo

HUR Hold-up ropes

HV High voltage

IH Inverted hierarchy

ITR In-Time-Ratio

LAB Linear alkylbenzene

LS Liquid scintillator

LY Light yield

MC Monte Carlo simulation

MCMC Markov Chain Monte Carlo

MTC/A Master Trigger Card/Analog

MTC/D Master Trigger Card/Digital

MR Molar ratio

NH Normal hierarchy

NHits Number of triggered (hit) PMTs in an event

NMR Nuclear magnetic resonance

PDF Probability distribution function

PMNS matrix Pontecorvo-Maki-Nakagawa-Sakata matrix

PMT Photomultiplier tube

PMTIC PMT Interface Card

PPO 2,5-Diphenyloxazole

PSD Pulse shape discrimination

PSUP PMT support structure

RAT Reactor Analysis Tool

ROI Region of interest

SM Standard Model of Particle Physics

TAC Time to Analog Converter

TeA Telluric acid

TeBD complex Tellurium butanediol complex

UPW Ultra pure water

Appendix B


List of Amines

Before DDA was selected for use in SNO+, a number of other amines were tested as potential stabilisers. The list below describes these compounds along with the tests performed and reasons for rejecting them.


Hexadecylamine (HDA) and Tetradecylamine (TDA)

Description: Primary amines with linear hydrocarbon chains of 16 and 14 carbons respectively. White powders at room temperature.

Chemical structures:



Hexadecylamine



Tetradecylamine

LY improvement: Yes. 45 – 50% for 1:1 amine:Te MR.

Tests performed: Duplicate samples for both amines with 1:1 amine:Te MR added to 0.5% Te solution. 48 hour humidity exposure.




Figure B.1: Photo of precipitation formed in a sample with tetradecylamine after humidity exposure.

Reason for rejection: Precipitation formed at the bottom of the vials within 5 days post humidity exposure (see Figure B.1). This timescale is too short for this to be a sign of a crash. Controls (samples with no amine added) in the same test did not form this precipitation. Furthermore, LY is stable over time, which suggests solubility issues of the amines rather than sample crashing.

Tris(2-ethylhexyl)amine (TEHA)

Description: Branched non-linear tertiary amine. Transparent liquid at room temperature.

Chemical structure:



LY improvement: No. Comparable $\sim 10\%$ decrease in LY for all 3 tested concentrations.

Tests performed: 48 hour humidity exposure test with concentrations corresponding to 0.0, 0.5, 1.0 and 1.4 TEHA:Te MR. Tested with 1% Te samples.

Reason for rejection: No signs of protection against water attack combined with decrease in LY. Samples with TEHA developed flakes after approximately 1 month after humidity exposure, i.e. on the same time frame as unprotected samples undergoing the same exposure. Curves of LY as a function of time also follow the trend of unprotected samples and show signs of instability after a few days.

Cocoamine

Description: Transparent liquid at room temperature with molecular formula $C_{38}H_{46}N_2O_8$.

Chemical structure: Various forms.

LY improvement: Yes. Up to 50% higher LY for 1%Te samples with 0.25 cocoamine:Te MR.


Tests performed: 48 hour humidity exposure with 0.1, 0.25, 0.5, 1 and 2 cocoamine:Te MR. Showed good protection against water attack from LY monitoring during the first week after exposure.

Reason for rejection: Obtained from biological sources. Large $^{14}C/^{12}C$ ratio is unacceptable radioactive contamination.

N,N-Dimethyltetradecylamine (DMTDA)

Description: Tertiary amine with linear hydrocarbon chains of 14 carbons. Clear liquid at room temperature.

Chemical structure:



LY improvement: Yes. 5 – 10% increase for 1:1 DMTDA:Te MR.


Tests performed: Duplicate samples with 1:1 DMTDA:Te MR added to 0.5% Te solution. 48 hour humidity exposure. Evidence of water attack protection.

Reason for rejection: Decreased transparency of samples after humidity exposure, improved at lower temperature environment. Worse LY than with other suitable candidates.

N,N-Dimethyldecylamine (DMDA)

Description: Tertiary amine with linear hydrocarbon chains of 10 carbons respectively. Clear liquid at room temperature.

Chemical structure:



LY improvement: Yes. ~ 20% LY increase for 0.25 and 0.5 DMDA:Te MR.

Tests performed: 48 hour humidity exposure with 0, 0.1, 0.25, 0.5, 1 and 2 DMDA:Te MR. Showed good protection against water attack from both LY and visual monitoring for over 2 years post exposure.

Reason for rejection: Flash point of 92°C [122] is lower than for DDA (113°C), otherwise DMDA appears as suitable candidate. Given humidity exposure tests yield similar results, DDA is selected.

Appendix C

^{60}Co and ^{208}Tl Decay Schemes




Figure C.1: ^{60}Co decay scheme. Green diagonal lines represent β decays and their relative probabilities. Horizontal lines represent γ rays, with dashed lines for very rare de-excitations and orange for de-excitations with probability $> 10\%$. The relative γ intensity is also explicitly given by purple numbers. Data taken from [64].




Figure C.2: ^{208}Tl decay scheme. Green diagonal lines represent β decays and their relative probabilities. Horizontal lines represent γ rays, with dashed black lines for very rare de-excitations, full black lines for de-excitations with probability between 0.1 and 10 % and orange for de-excitations with probability > 10 %. The relative γ intensity is also explicitly given by purple numbers. Data taken from [64].

Appendix D

Correcting for Detector State in Energy Reconstruction

Introduction

The parameters used in `EnergyRThetaFunctional` to convert from N_{Hits} in an event to energy are tuned in a process referred to as fitter coordination, which uses Monte Carlo simulation of particles with known position and energy. This simulation is created using a certain detector state, which specifies the number of PMTs on high voltage, their channel thresholds and whether they are properly calibrated. However, in reality the detector state is not fixed and can be easily modified. In particular, PMTs can be temporarily turned off due to high voltage problems and, hence, the number of online PMTs can change multiple times per day. The detector operation is split into "runs", which typically last 1 hour. The detector state will remain fixed during that period but can change from one run to another. To account for the different setting in runs and fitter coordination, a detector state correction was added to the energy reconstruction algorithm and is described below.

Method

The energy estimator H is obtained for each event according to equation 5.3. It is also possible to write the expression as

$$H = - \sum_{i=1}^S N_{PMT}^i \cdot \mu^i, \quad (\text{D.1})$$

where μ^i is the Poisson mean of hits per PMT in segment i . This is distinct from the observed N_{Hits}^i , which gives the number of triggered PMTs without distinguishing whether the trigger was caused by one or more photoelectrons. In turn, the estimator relates to the energy as $E = \alpha H$, where α is the relevant constant obtained during fitter coordination and saved in the RAT database.

When the detector state changes compared to what was used in coordination, the differences in two variables must be corrected: the number of active PMTs and the channel efficiency. The former gives the number of inward looking PMTs that have HV on, have been calibrated properly and are flagged as functional in the database entry for the given run. Assuming PMTs break homogeneously throughout the detector, the difference in active PMTs in a segment i changes as follows:

$$N_{PMT}^i \rightarrow N'^i_{PMT} = f N_{PMT}^i \quad f = \frac{\# \text{ of active PMTs in run}}{\# \text{ of active PMTs in coordination}}. \quad (\text{D.2})$$

On the other hand, the difference in channel efficiency modifies how likely is a photoelectron to be detected and, hence, effectively scales μ^i :

$$\mu^i \rightarrow \mu'^i = g \mu^i \quad g = \frac{\text{mean channel efficiency in run}}{\text{mean channel efficiency in coordination}}. \quad (\text{D.3})$$

For simplicity, only the average channel efficiency is compared here. Therefore, an assumption is made again that the channel thresholds change approximately by the same amount in each segment. This is generally true as the thresholds

are usually adjusted globally.

Combining equations D.1 to D.3, the energy estimator calculated for a given event is

$$H' = - \sum_{i=1}^S N_{PMT}^i \cdot \mu^i = -gf \sum_{i=1}^S N_{PMT}^i \cdot \mu^i = gfH. \quad (D.4)$$

Therefore, it is possible to obtain a corrected constant of proportionality α' that converts H' into energy, simply by scaling the parameter α from coordination:

$$E = \alpha H = \alpha' H' = \alpha' gfH \quad \implies \quad \alpha' = \frac{\alpha}{gf}. \quad (D.5)$$

`EnergyRThetaFunctional` has now been modified to calculate gf for each run and apply it to α . In order to do that, the fitter coordination now saves the number of active PMTs and average channel efficiency used in the processes into the database along with the tuned parameters. The corresponding quantities for the given run are obtained during reconstruction using existing functionality in RAT.

Performance

To test the performance of the correction, electrons with a range of energies between 1 and 10 MeV were simulated with variable numbers of the PMTs switched off at random. The fractional bias in the reconstructed energy with the detector state correction is plotted in Figure D.1. In all cases, the fitter was coordinated with all PMTs online. It is shown that the bias remains $< 1\%$ even with almost half the PMTs turned off. The energy of electrons simulated with all PMTs switched on is plotted in purple to show the case when detector state in coordination and run are matched.

For comparison, the energy bias of electrons simulated with no correction and "best realistic" detector state, which only excludes ~ 500 PMTs that have been deemed permanently damaged, is plotted in orange. Even this best achievable




Figure D.1: Comparison of reconstructed energy bias with (blues and greens) and without (orange) the detector state correction. Simulation with all PMTs turned on as in coordination shown in purple. Each point represents the mean of a Gaussian fit to a histogram of $(E_{reconstructed} - E_{true})/E_{true}$, where $E_{reconstructed}$ and E_{true} are the reconstructed and simulated energy, respectively. Vertical error bars arise from errors of the Gaussian fit.

state introduces a bias of $\sim 2\%$ when the offline PMTs are not accounted for.

The implementation of the correction assumes that the detector state changes equally in each detector segment that is used in the calculation of the energy estimator. To examine how prone is the correction to deviations from this assumption, simulations were made with one or more neighbouring electronic crates turned off. One crate controls ~ 500 adjacent PMTs placed in a half-circular segment along the Z axis. As shown in Figure D.2, the energy bias remains $< 1\%$ with one entire crate off and only increases significantly as the number of disabled crates rises. During stable data taking, runs with one or more crates off are rare and are currently being excluded from data analysis. Therefore, all runs used for analysis are expected to have more homogeneous distribution of inactive PMTs and, hence, the energy correction will be valid for these runs.




Figure D.2: Reconstructed energy bias with detector state correction and areas of adjacent PMTs turned off (blues and greens). Bias without the correction and the best realistic state (orange) and all PMTs online (purple) also shown. Each point represents the mean of a Gaussian fit to a histogram of $(E_{reconstructed} - E_{true})/E_{true}$, where $E_{reconstructed}$ and E_{true} are the reconstructed and simulated energy, respectively. Vertical error bars arise from errors of the Gaussian fit.

Summary

To rectify the energy bias caused by different detector states in fitter coordination and analysed data or simulation, a correction was implemented in the `EnergyRThetaFunctional` reconstruction algorithm. Corrections for both the number of active PMTs and average channel efficiency were applied as linear scaling. The validity of this addition was tested with both individual PMTs turned off at random and electronics crates disabled. With the correction, the energy bias remains sub-percent level with almost half of the PMTs or one full crate inactive. This is compared to $\sim 2\%$ bias with no correction and only 500 PMTs inactive. The detector state is expected to continue to change once scintillator data taking commences and, therefore, the ability of the `EnergyRThetaFunctional` method to produce valid results for various detector states without the need of coordination will be crucial.

Appendix E

Tuning Classifier Weights on Different Backgrounds

This appendix discusses the effect of tuning the timing and topological classifiers, described in section 5.3, on different types backgrounds. Unfortunately, at the time when the official simulation used in this thesis (RAT6.16.3) was being produced, the classifiers were only assumed to serve for rejection of external backgrounds. Therefore, all the weights were calculated with the most dangerous external background for the $0\nu\beta\beta$ search: ^{208}Tl .

However, there are multiple locations in the detector where this isotope can decay: the AV wall, the PMT glass, the water surrounding the AV, the dust on AV surfaces and the hold-down and hold-up ropes. Monte Carlo simulations of four of these sources were used to calculate $\Delta \log \mathcal{L}$ weights for both the timing and topological classifiers. All eight resulting classifiers were then run over independent Monte Carlo samples in order to assess the possibility of using weights tuned with one background source to discriminate between signal and background events of another type. Figures E.1 and E.2 display the results for the timing and topological classifiers, respectively.

These plots show the fraction of external ^{208}Tl event of a given type passing

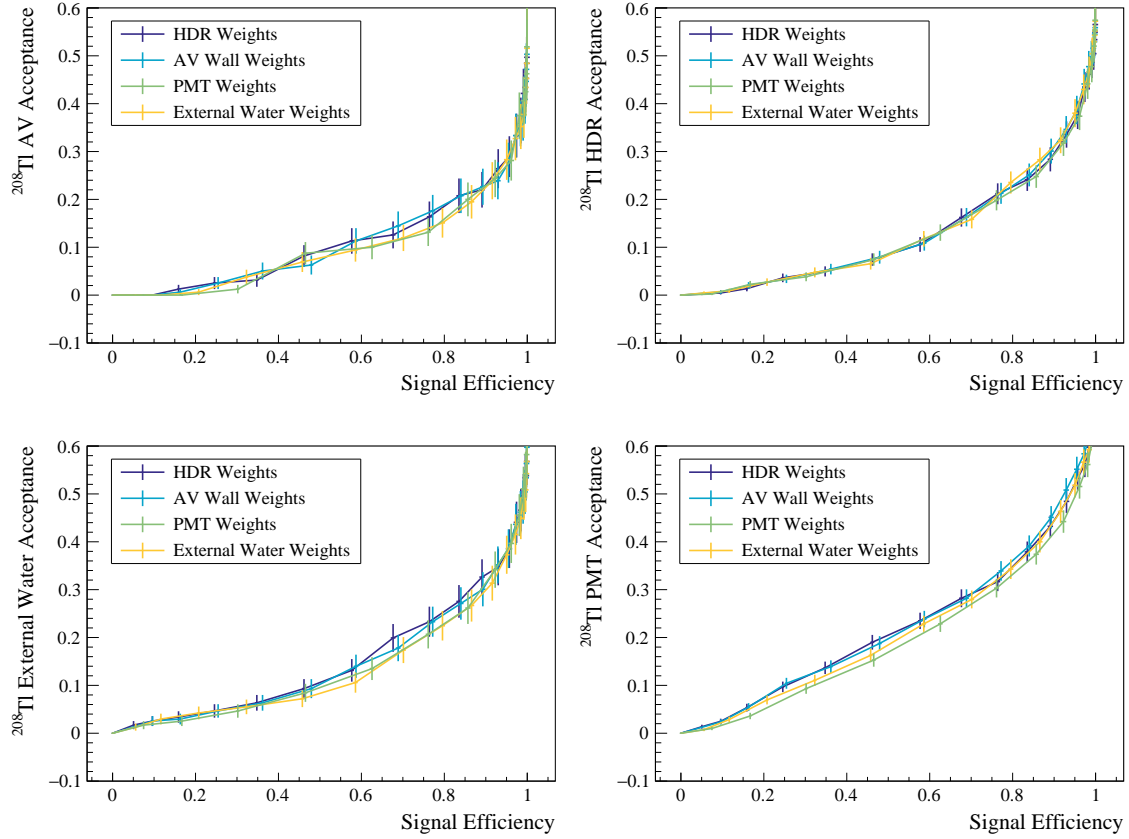


Figure E.1: Fraction of external ^{208}Tl events from various sources passing the timing classifier cut as a function of fraction of accepted $0\nu\beta\beta$ events. Curves shown for weights tuned on MC simulation of ^{208}Tl from hold-down ropes (purple), AV wall (light blue), PMT glass (green) and water surrounding the AV (orange).

the classifier cut as function of signal efficiency. In each plot, one of the curves was obtained with classifier tuned on the corresponding background source (i.e hold-down rope events classified with hold-down rope weights), whereas weights for different backgrounds were used for the remaining three curves.

In all of the plots, negligible difference in the background rejection is seen between the various weights. Therefore, it is concluded that weights tuned on one external ^{208}Tl source can be used for rejection of other external ^{208}Tl backgrounds with minimal performance loss. This is due to the shared features in time residual and angular distributions of the multi-site backgrounds described in section 5.3.1 and justifies the usage of only ^{208}Tl AV wall weights throughout this thesis.

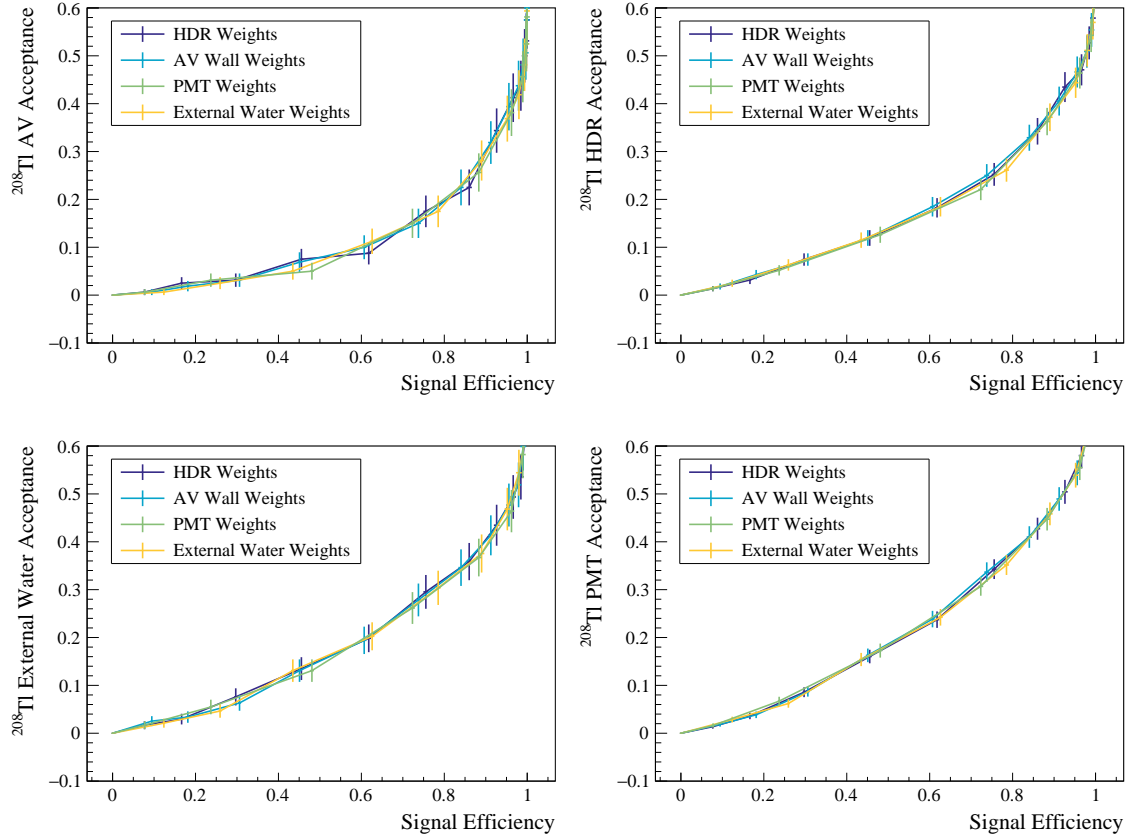


Figure E.2: Fraction of external ^{208}Tl events from various sources passing the topological classifier cut as a function of fraction of accepted $0\nu\beta\beta$ events. Curves shown for weights tuned on MC simulation of ^{208}Tl from hold-down ropes (purple), AV wall (light blue), PMT glass (green) and water surrounding the AV (orange).

It was also demonstrated that ^{208}Tl weights can be used to cut external ^{214}Bi backgrounds, which is important at lower energies for solar neutrino analyses in unloaded scintillator phase [123]. In addition, the classifiers have been used for internal background rejection, with the most notable mention of ^{60}Co . Results in section 5.3.1 as well as Chapter 6 show that the timing classifier is valuable in mitigation of internal multi-site events despite being tuned on external backgrounds. However, it might be possible to improve on these results even further if specific tuning was available for internal backgrounds.

The execution of the classifier algorithm takes 0.3 ms per call, which is 3 orders of magnitude faster than the position and energy reconstruction of the

event. Therefore, it is not problematic to run the classifiers multiple times for each event. As a result, the classifiers are run with all of the available weights as a part of the `ScintFitter` reconstruction routine, despite the fact that there is minimal difference in the classifier performance for the different weights.

Appendix F

Event Rates and Fit Parameters

Table F.1: Event types included in the likelihood analysis discussed in Chapter 6. Expected background rates taken from [88] are listed in second column. Third column gives the fraction of events of a given type passing the analysis cuts from section 6.5. Whether the normalisation of a given event types is floated in the likelihood fit is indicated in the last column. Event types are listed in order of descending count contribution to the extended ROI after cuts. For external backgrounds, their source is stated after the isotope name with the following acronyms: AV = acrylic vessel, HDR = hold-down ropes, HUR = hold-up ropes, IR = internal (calibration) ropes, EW = external water, OD = outer dust on AV, ID = inner dust on AV. Rate of cosmogenic backgrounds was calculated as a five year average after a one year cool down underground and purification.

Event Type	Expected Yearly Rate in Full Detector	Fraction of Counts Passing Cuts [%]	Fit Parameter
$0\nu\beta\beta$	-	73.90	Yes
$2\nu\beta\beta$	5.41×10^6	0.96	Yes
^{214}Bi AV	1.28×10^7	0.18	Yes
^{208}Tl AV	1.50×10^6	1.01	Yes
^{210}Tl HDR	1.55×10^6	0.78	Yes
^{214}Bi EW	4.62×10^7	0.02	Yes
PMT $\beta\gamma$ ¹	7.18×10^4	9.16	Yes
^{208}Tl OD	4.60×10^5	0.98	Yes

continued on next page

Event Type	Expected Yearly Rate in Full Detector	Fraction of Counts Passing Cuts [%]	Fit Parameter
^{208}Tl EW	1.96×10^6	0.16	Yes
^{234m}Pa	3.95×10^5	0.71	Yes
^{214}Bi HDR	2.72×10^6	0.08	Yes
^{228}Ac	5.57×10^4	3.48	Yes
^{210}Tl HUR	1.67×10^5	0.68	Yes
^{214}Bi OD	7.75×10^5	0.14	Yes
^{214}Bi IR	4.97×10^3	9.38	Yes
^{214}Bi HUR	2.92×10^5	0.07	Yes
$^{214}\text{BiPo}$	3.95×10^5	0.04	Yes
^8B solar ν_e	9.54×10^2	12.68	Yes
$^{212}\text{BiPo}$	5.57×10^4	0.21	Yes
^{208}Tl ID	2.48×10^4	0.30	Yes
(α, n) ^{13}C LS ²	3.01×10^2	21.81	Yes
^{214}Bi ID	4.15×10^4	0.11	Yes
^8B solar ν_μ	3.15×10^2	13.22	Yes
^{208}Tl	2.01×10^4	0.14	Yes
(α, n) ^{13}C AVin LS ²	8.98×10^2	1.90	Yes
(α, n) ^{13}C AVin AV ²	6.74×10^2	2.02	Yes
^{210}Tl	8.30×10^1	15.56	Yes
(α, n) ^{13}C AVout AV ²	6.33×10^2	1.56	Yes
^{130}I	5.47	68.50	Yes
^{130m}I	6.50	3.54	Yes
^{102m}Rh	2.71×10^{-1}	46.30	No
^{124}Sb	2.29×10^{-2}	53.01	No
^{88}Y	1.89×10^{-2}	63.21	Yes
^{110m}Ag	9.11×10^{-3}	69.20	No

continued on next page

Event Type	Expected Yearly Rate in Full Detector	Fraction of Counts Passing Cuts [%]	Fit Parameter
^{22}Na	5.77×10^{-3}	62.50	Yes
^{102}Rh	7.51×10^{-2}	4.47	No
^{60}Co	3.46×10^{-3}	67.30	Yes
^{46}Sc	1.90×10^{-3}	66.33	No
^{42}K	3.44×10^{-3}	31.20	No
^{68}Ga	3.84×10^{-3}	24.60	No
^{106}Rh	9.40×10^{-4}	27.60	No
^{44}Sc	1.55×10^{-4}	60.60	No
^{126m}Sb	8.00×10^{-5}	60.61	No
^{126}Sb	5.60×10^{-5}	46.90	No
^{110}Ag	1.23×10^{-4}	12.90	No
^{58}Co	1.41×10^{-5}	5.29	No
^{82}Rb	1.35×10^{-6}	40.11	No
^{84}Rb	4.10×10^{-6}	10.90	No
^{26}Al	7.72×10^{-7}	23.01	No
^{56}Co	1.89×10^{-6}	8.72	No

¹PMT $\beta\gamma$ includes the weighted external backgrounds from the PMTs, with the energy and radial distributions of γ s and β s generated based on the model developed by W. Heintzelman [124].

²For (α, n) interactions on ^{13}C , the α can be either produced inside the liquid scintillator volume (LS) or on one of the surfaces of the AV (AVin and AVout). In the latter case, the direction of the α dictates whether it then interacts with a ^{13}C inside the scintillator (LS) or the acrylic (AV). This gives the total of 4 categories of this background modelled separately.

References

- [1] SNO Collaboration, *Direct evidence for neutrino flavor transformation from neutral current interactions in the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **89** (2002) 011301.
- [2] Super-Kamiokande Collaboration, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562–1567.
- [3] M. Fukugita and T. Yanagida, *Baryogenesis Without Grand Unification*, Phys. Lett. B **174** (1986) 45–47.
- [4] M. Gell-Mann, P. Ramond, and R. Slansky, *Complex Spinors and Unified Theories*, Conf. Proc. C **790927** (1979) 315–321.
- [5] G. Lazarides, Q. Shafi, and C. Wetterich, *Proton Lifetime and Fermion Masses in an $SO(10)$ Model*, Nucl. Phys. B **181** (1981) 287–300.
- [6] C. Cowan et al., *Detection of the free neutrino: A Confirmation*, Science **124** (1956) 103–104.
- [7] M. Tanabashi et al. (Particle Data Group), *Review of Particle Physics*, Phys. Rev. D **98** (2018) 030001.
- [8] C. Giunti and C. W. Kim, *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, 2007.
- [9] M. Tanabashi et al. (Particle Data Group), *Neutrino Masses, Mixing, and Oscillations*, Phys. Rev. D **98** (2018) 030001.
- [10] SNO Collaboration, *Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory*, Phys. Rev. C **88** (2013) 025501.
- [11] J. Hewett et al., *Fundamental physics at the intensity frontier*, arXiv:1205.2671 [hep-ex].
- [12] KATRIN Collaboration, *Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN*, Phys. Rev. Lett. **123** (2019) no. 22, 221802.
- [13] Planck Collaboration, *Planck 2018 results. VI. Cosmological parameters*, arXiv:1807.06209 [astro-ph.CO].

- [14] E. Majorana, *Symmetric theory of electrons and positrons*, *Il Nuovo Cimento* **14** (1937) 201–233.
- [15] R. Mohapatra and A. Smirnov, *Neutrino Mass and New Physics*, *Ann. Rev. Nucl. Part. Sci.* **56** (2006) 569–628.
- [16] W. Furry, *On transition probabilities in double beta-disintegration*, *Phys. Rev.* **56** (1939) 1184–1193.
- [17] M. Goepfert-Mayer, *Double beta-disintegration*, *Phys. Rev.* **48** (1935) 512–516.
- [18] P. Bamert, C. Burgess, and R. Mohapatra, *Multi-majoron modes for neutrinoless double-beta decay*, *Nucl. Phys. B* **449** (1995) 25 – 48.
- [19] M. Hirsch et al., *On the observability of Majoron emitting double beta decays*, *Phys. Lett. B* **372** (1996) 8 – 14.
- [20] F. T. Avignone III et al., *Double Beta Decay, Majorana Neutrinos, and Neutrino Mass*, *Rev. Mod. Phys.* **80** (2008) 481–516.
- [21] KamLAND-Zen Collaboration, *Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen*, *Phys. Rev. Lett.* **117** (2016) no. 8, 082503.
- [22] P. De Salas et al., *Neutrino Mass Ordering from Oscillations and Beyond: 2018 Status and Future Prospects*, *Front. Astron. Space Sci.* **5** (2018) 36.
- [23] S. Sarkar, *Measuring the baryon content of the universe: BBN vs CMB*, arXiv:0205116 [astro-ph].
- [24] P. Di Bari, *An introduction to leptogenesis and neutrino properties*, *Contemp. Phys.* **53** (2012) no. 4, 315–338.
- [25] S. Umehara et al., *Neutrino-less double-beta decay of Ca-48 studied by CaF₂(Eu) scintillators*, *Physical Review C* **78** (2008) 058501.
- [26] GERDA Collaboration, *Probing Majorana neutrinos with double- β decay*, *Science* **365** (2019) 1445–1448.
- [27] NEMO-3 Collaboration, *Investigation of double-beta decay with the NEMO-3 detector*, *Phys. At. Nucl.* **74** (2011) no. 2, 312–317.
- [28] NEMO-3 Collaboration and SuperNEMO Collaboration, *Latest results of NEMO-3 experiment and present status of SuperNEMO*, *Nucl. Part. Phys. Proc.* **273-275** (2016) 1765–1770.
- [29] NEMO-3 Collaboration, *Measurement of the $2\nu\beta\beta$ Decay Half-Life and Search for the $0\nu\beta\beta$ Decay of ^{116}Cd with the NEMO-3 Detector*, *Phys. Rev. D* **95** (2017) 012007.
- [30] CUORE Collaboration, *Improved Limit on Neutrinoless Double-Beta Decay in ^{130}Te with CUORE*, *Phys. Rev. Lett.* **124** (2020) no. 12, 122501.

- [31] NEMO-3 Collaboration, *Measurement of the $2\nu\beta\beta$ decay half-life of ^{150}Nd and a search for $0\nu\beta\beta$ decay processes with the full exposure from the NEMO-3 detector*, Phys. Rev. D **94** (2016) 072003.
- [32] IGEX Collaboration, *IGEX ^{76}Ge neutrinoless double-beta decay experiment: Prospects for next generation experiments*, Phys. Rev. D **65** (2002) 092007.
- [33] H. Klapdor-Kleingrothaus et al., *Latest results from the Heidelberg-Moscow double beta decay experiment*, Eur. Phys. J. A **12** (2001) 147–154.
- [34] GERDA Collaboration, *The GERDA experiment for the search of $0\nu\beta\beta$ decay in ^{76}Ge* , Eur. Phys. J. C **73** (2013) no. 3, 2330.
- [35] GERDA Collaboration, *Background-free search for neutrinoless double- β decay of ^{76}Ge with GERDA*, Nature **544** (2017) no. 7648, 47–52.
- [36] Majorana Collaboration, *The Majorana Demonstrator Status and Preliminary Results*, EPJ Web Conf. **178** (2018) 01006.
- [37] LEGEND Collaboration, *The Large Enriched Germanium Experiment for Neutrinoless Double Beta Decay (LEGEND)*, AIP Conf. Proc. **1894** (2017) no. 1, 020027.
- [38] K. Zuber, *COBRA - Double beta decay searches using CdTe detectors*, Phys.Lett. B **519** (2018) no. 1-2, 1–7.
- [39] D. R. Artusa et al., *Exploring the neutrinoless double beta decay in the inverted neutrino hierarchy with bolometric detectors*, Eur. Phys. J. C **74** (2014) no. 3096, 1–19.
- [40] C. Brofferio, O. Cremonesi, and S. Dell’Oro, *Neutrinoless Double Beta Decay Experiments With TeO_2 Low-Temperature Detectors*, Front. in Phys. **7** (2019) 86.
- [41] CUORE Collaboration, *Results from the Cuore Experiment*, Universe **5** (2019) no. 10, 1–10.
- [42] CUPID Collaboration, *R&D towards CUPID (CUORE Upgrade with Particle Identification)*, arXiv:1504.03612 [physics.ins-det].
- [43] R. Arnold et al., *Technical design and performance of the NEMO 3 detector*, Nucl. Instrum. Meth. A **536** (2005) 79–122.
- [44] SuperNEMO Collaboration, *Probing New Physics Models of Neutrinoless Double Beta Decay with SuperNEMO*, Eur. Phys. J. C **70** (2010) 927–943.
- [45] NEMO Collaboration, *First results of the search of neutrinoless double beta decay with the NEMO 3 detector*, Phys. Rev. Lett. **95** (2005) 182302.
- [46] NEMO-3 Collaboration, *Measurement of the Double Beta Decay Half-life of ^{130}Te with the NEMO-3 Detector*, Phys. Rev. Lett. **107** (2011) 062504.

- [47] NEMO-3 Collaboration, *Measurement of the two neutrino double beta decay half-life of Zr-96 with the NEMO-3 detector*, Nucl. Phys. A **847** (2010) 168–179.
- [48] C. Augier et al., *Technical design and performance of the NEMO 3 detector*, Nucl. Instrum. Meth. A **536** (2005) 79–122.
- [49] H. J. Hilke, *Time projection chambers*, Rep. Prog. Phys. **73** (2010) no. 11, 116201.
- [50] M. Auger et al., *The EXO-200 detector, part I: Detector design and construction*, JINST **7** (2012) P05010.
- [51] EXO-200 Collaboration, *Search for Neutrinoless Double-Beta Decay with the Upgraded EXO-200 Detector*, Phys. Rev. Lett. **120** (2018) no. 7, 2–7.
- [52] nEXO Collaboration, *nEXO Pre-Conceptual Design Report*, arXiv:1805.11142 [physics.ins-det].
- [53] nEXO Collaboration, *Imaging individual barium atoms in solid xenon for barium tagging in nEXO*, Nature **569** (2019) no. 7755, 203–207.
- [54] A. McDonald et al., *Demonstration of Single Barium Ion Sensitivity for Neutrinoless Double Beta Decay using Single Molecule Fluorescence Imaging*, Phys. Rev. Lett. **120** (2018) no. 13, 132504.
- [55] BOREXINO Collaboration, *Comprehensive measurement of pp-chain solar neutrinos*, Nature **562** (2018) no. 7728, 505–510.
- [56] BOREXINO Collaboration, *Light propagation in a large volume liquid scintillator*, Nucl. Instrum. Meth. A **440** (2000) no. 2, 360–371.
- [57] H. W. C. Tseung, J. Kaspar, and N. Tolich, *Measurement of the dependence of the light yields of linear alkylbenzene-based and EJ-301 scintillators on electron energy*, Nucl. Instrum. Meth. A **654** (2011) 318–323.
- [58] KamLAND-Zen Collaboration, *Measurement of the double- β decay half-life of ^{136}Xe with the KamLAND-Zen experiment*, Phys. Rev. C **85** (2012) 045504.
- [59] R. Nakamura et al., *Research and development toward KamLAND2-Zen*, J. Phys. Conf. Ser. **1468** (2020) 012256.
- [60] SNO+ Collaboration, *Current Status and Future Prospects of the SNO+ Experiment*, Adv. High Energy Phys **2016** (2016) no. 6194250, 1–21.
- [61] SNO Collaboration, *The Sudbury neutrino observatory*, Nucl. Instrum. Meth. A **449** (2000) 172–207.
- [62] S. D. Biller, *Probing Majorana neutrinos in the regime of the normal mass hierarchy*, Phys. Rev. D **87** (2013) 1–4.
- [63] M. Berglund and M. Wieser, *Isotopic compositions of the elements 2009 (IUPAC technical report)*, Pure Appl. Chem. **83** (01, 2011) 397–410.

- [64] National Nuclear Data Center, *NuDat 2.8*, <https://www.nndc.bnl.gov/nudat2/>.
- [65] SNO+ Collaboration, *The SNO+ Experiment*, paper in preparation .
- [66] SNO Collaboration, *Cosmogenic Neutron Production at the Sudbury Neutrino Observatory*, *Phys. Rev. D* **100** (2019) 112005.
- [67] S. Biller et al., *Measurements of photomultiplier single photon counting efficiency for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **432** (1999) 364–373.
- [68] G. Doucas et al., *Light concentrators for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **370** (1996) 579–596.
- [69] SNO Collaboration, *Sudbury neutrino observatory*, *Nucl. Instrum. Meth. A* **449** (2000) 172–207.
- [70] R. N. Mohapatra and A. Pérez-Lorezana, *Neutrino mass, proton decay, and dark matter in TeV scale universal extra dimension models*, *Phys. Rev. D* **67** (2003) 075015.
- [71] D. Barducci, M. Fabbrichesi, and E. Gabrielli, *Neutral Hadrons Disappearing into the Darkness*, *Phys. Rev. D* **98** (2018) 035049.
- [72] X. He and S. Pakvasa, *Unparticle Induced Baryon Number Violating Nucleon Decays*, *Phys. Lett. B* **662** (2008) 259–263.
- [73] H. Ejiri, *Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei*, *Phys. Rev. C* **48** (1993) 1442–1444.
- [74] SNO+ Collaboration, *Search for invisible modes of nucleon decay in water with the SNO+ detector*, *Phys. Rev. D* **99** (2019) 032008.
- [75] SNO+ Collaboration, *Measurement of the ^8B solar neutrino flux in SNO+ with very low backgrounds*, *Phys. Rev. D* **99** (2019) 012012.
- [76] Super-Kamiokande Collaboration, *Solar Neutrino Measurements in Super-Kamiokande-IV*, *Phys. Rev. D* **94** (2016) 052010.
- [77] I. T. Coulter, *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, University of Oxford, 2013.
- [78] SNO+ Collaboration, *AV pictures during fill*, SNO+ DocDB **5839-v2** (2019) .
- [79] B. A. Moffat et al., *Optical calibration hardware for the Sudbury Neutrino Observatory*, *Nucl. Instrum. Meth. A* **554** (2005) 255–265.
- [80] B. A. Moffat, *The Optical Calibration of the Sudbury Neutrino Observatory*. PhD thesis, Queen’s University, 2001.

- [81] SNO+ Collaboration, *The calibration system for the photomultiplier array of the SNO+ experiment*, JINST **10** (2015) no. 03, P03002.
- [82] Geant4 Collaboration, *GEANT4 - A simulation toolkit*, Nucl. Instrum. Meth. A **506** (2003) no. 3, 250–303.
- [83] G. Horton-Smith, *GLG4sim page*, <https://www.phys.ksu.edu/personal/gahs/GLG4sim/>.
- [84] J. B. Birks, *The Theory and Practice of Scintillation Counting*. Pergamon Press Ltd., Oxford, 1964.
- [85] W. R. Leo, *Techniques for Nuclear and Particle Physics Experiments: A How-to Approach*. Springer-Verlag, 1994.
- [86] T. Kaptanoglu et al., *Optics Overview and Proposed Changes to RAT*, SNO+ DocDB **3461** (2016) .
- [87] L. Segui, *Scintillator Model : comparison between new data and old model and its performance in RAT*, SNO+ DocDB **2774** (2015) .
- [88] V. Lozza, *Expected radioactive backgrounds in SNO+*, SNO+ DocDB **507-v42** (2019) .
- [89] S. Biller et al., *A New Technique to Load ^{130}Te in Liquid Scintillator for Neutrinoless Double Beta Decay Experiments*, J. Phys. Conf. Ser. **888** (2017) 012084.
- [90] SNO+ Collaboration, *Tellurium loading into liquid scintillator*, paper in preparation .
- [91] J. J. Lidgard. Personal communication.
- [92] T. Kaptanoglu and K. Galban, *Light Yields With DDA*, SNO+ DocDB **5124-v2** (2018) .
- [93] S. Biller, *Overview of Chemistry Loading Development*, SNO+ DocDB **5732** (2019) .
- [94] National Center for Biotechnology Information, *N-Octylamine*, PubChem Database , <https://pubchem.ncbi.nlm.nih.gov/compound/Octylamine>.
- [95] National Center for Biotechnology Information, *N,N-Dimethyldodecylamine*, PubChem Database , https://pubchem.ncbi.nlm.nih.gov/compound/N_N-Dimethyldodecylamine.
- [96] S. Hans et al., *Purification of telluric acid for SNO+ neutrinoless double-beta decay search*, Nucl. Instrum. Meth. A **795** (2015) 132–139.
- [97] H. O’Keeffe, E. O’Sullivan, and M. Chen, *Scintillation decay time and pulse shape discrimination in oxygenated and deoxygenated solutions of linear alkylbenzene for the SNO+ experiment*, Nucl. Instrum. Meth. A **640** (2011) 119–122.

- [98] SNO Collaboration, *Measurement of the Cosmic Ray and Neutrino-Induced Muon Flux at the Sudbury Neutrino Observatory*, Phys. Rev. D **80** (2009) 012001.
- [99] M. J. D. Powell, *An efficient method for finding the minimum of a function of several variables without calculating derivatives*, Comput J. **7** (1964) no. 2, 155–162.
- [100] M. Mottram, *Updated functional form energy reconstruction : EnergyRThetaFunctional*, SNO+ DocDB **3488** (2015) .
- [101] J. Dunger, *Topological and Time Based Event Classification for Neutrinoless Double Beta Decay in Liquid Scintillator*. PhD thesis, University of Oxford, 2018.
- [102] J. Dunger and S. D. Biller, *Multi-site Event Discrimination in Large Liquid Scintillation Detectors*, Nucl. Instrum. Meth. A **943** (2019) 162420.
- [103] G. Cowan, *Statistical Data Analysis*. Oxford University Press Inc., 1998.
- [104] R. A. Fisher, *The Use of Multiple Measurements in Taxonomic Problems*, Annals of Eugenics **7** (1936) no. 2, 179–188.
- [105] J. Bahcall, *Software and data for solar neutrino research*, 2005. <http://www.sns.ias.edu/~jnb/>.
- [106] BOREXINO Collaboration, *Improved measurement of ^8B solar neutrinos with 1.5kt \cdot y of Borexino exposure*, Phys. Rev. D **101** (2020) 062001.
- [107] E. W. Beier, *Study of Direction Reconstruction in Scintillator*, SNO+ DocDB **3503** (2016) .
- [108] M. Mottram, *Cherenkov reconstruction in scintillator at 2.5 MeV refracted light path*, SNO+ DocDB **3116** (2015) .
- [109] KamLAND Collaboration, *First Results from KamLAND: Evidence for Reactor Antineutrino Disappearance*, Phys. Rev. Lett. **90** (2003) no. 2, 6.
- [110] J. Wilson, *BiPo Rejection Factors Summary*, SNO+ DocDB **4183** (2017) .
- [111] K. Majumdar, *On the Measurement of Optical Scattering and Studies of Background Rejection in the SNO+ Detector*. PhD thesis, University of Oxford, 2015.
- [112] R. Lane, *Counting analysis update*, SNO+ DocDB (2017) .
- [113] T. Kaptanoglu, *Out-Of-Window BiPo Backgrounds For the $0\nu\beta\beta$ Phase*, SNO+ DocDB **5531** (2019) .
- [114] V. Lozza and J. Petzoldt, *Cosmogenic activation of a natural tellurium target*, Astropart. Phys. **61** (2015) 62–71.

- [115] H. Ejiri and S. Elliott, *Charged current neutrino cross section for solar neutrinos, and background to $\beta\beta(0\nu)$ experiments*, Phys. Rev. C **89** (2014) 055501.
- [116] I. Coulter, *Alpha-n in the Te Diol*, SNO+ DocDB **4332** .
- [117] R. M. Neal, *MCMC using hamiltonian dynamics*, Handbook of Markov Chain Monte Carlo (2011) 113–162.
- [118] G. Cowan et al., *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C **71** (2011) 1554.
- [119] J. Kotila and F. Iachello, *Phase space factors for double- β decay*, Phys. Rev. C **85** (2012) 034316.
- [120] J. Barea, J. Kotila, and F. Iachello, *Nuclear matrix elements for double- β decay*, Phys. Rev. C **87** (2013) 014315.
- [121] J. Wilson, *BiPo Rejection on Te-Diol*, SNO+ DocDB **4134** (2017) .
- [122] National Center for Biotechnology Information, *N,N-Dimethyldecylamine*, PubChem Database , https://pubchem.ncbi.nlm.nih.gov/compound/N_N-Dimethyldecylamine.
- [123] T. Kroupová, *External background classifiers*, SNO+ DocDB **5687-v3** (2019) .
- [124] W. Heintzelman, *Angular distribution of surviving γ s from PMT β - γ events*, SNO+ DocDB **1712** (2013) .