First Data Reconstruction and Inverse Beta Decay Analysis at the Large Scale SoLid Prototype Detector

Daniel Martin Saunders

A dissertation submitted to the University of Bristol in accordance
with the requirements for award of the degree of

Doctor of Philosophy

in the

Faculty of Science,

School of Physics.

February 2017

(Word count: 36200)

Abstract

The SoLid experiment is one of several new very short baseline neutrino experiments, searching for oscillations of anti-electron neutrinos into a new sterile flavour state. Motivations for such a search arise from various anomalies observed by previous experiments. This includes both the Reactor and Gallium anomalies, where many experiments have observed significant deficits in the rate of neutrinos at short baselines from reactors and radioactive sources. Critical to the success of the experiment is high energy resolution and effective background rejection. Using the Belgian Research Reactor 2 (BR2) nuclear reactor as a source, the experiment uses a novel detector technology that has been designed to combat many of the challenges of performing experiments very close to reactors. It is composed of around 13 thousand $5\,\mathrm{cm}^3$ composite scintillator cubes. The bulk of these cubes consists of the plastic scintillator Polyvynil-Toluene, with layers of ^{6}LiF:ZnS placed on the cube faces, allowing the detector to observe both neutrons and electromagnetic interactions. Neutrinos are detected via inverse beta decays (IBD), where neutrinos from the reactor are captured on protons from the detector volume. The high level of segmentation allows the topology of neutrino interactions to be studied to previously unseen precision, allowing for methods to distinguish the neutrino signal from the large backgrounds. The full 1.6 tonne detector will be commissioned in the spring of 2017, and will be placed just 6m from the reactor core. It aims to provide first results of the sterile search within one year.

A large scale 288 kg prototype using the technology was built and commissioned in spring 2015. This thesis describes the first data reconstruction and neutrino analysis based on the data from this prototype. The data recorded include extensive reactor off and source runs, as well as a short reactor on run. Reconstruction algorithms have been developed to find neutrino candidate interactions from raw detector data. To ensure high energy resolution, a technique using cosmic muons has been developed to calibrate the energy response of the cubes, resulting in a cube-to-cube variation of around 2%. Using these calibrations and event reconstruction algorithms, a full IBD analysis has been performed. It is found that the topological information of reconstructed IBD interactions, provided by the segmented design of the detector, can reduce the neutrino backgrounds by over two orders of magnitude after all other signal selection criteria have been applied. Comparisons of the reactor on and off data, using both traditional and machine learning techniques, show a statistically insignificant neutrino signal, in line with expectations given the observed background levels and the size of the dataset. These studies, backed up by studies on simulated data, validate the background estimation methods developed in this thesis. They also provide the blueprint for event reconstruction and event classification in the full detector. This work has informed many design aspects of the full SoLid detector, such as shielding requirements, trigger algorithms, and plans for the software infrastructure.

Acknowledgements

I am grateful to the Science and Technology Facilities Council for their funding through this research, and to the European Research Council for funding the SoLid experiment.

Many individuals also deserve gratitude.

I give enormous thanks to my supervisor Jonas Rademacker. This is the second time we have worked together, including my MSci degree, and I hope not the last! I am especially grateful for the freedom I was given to explore my own ideas, for quick feedback about possible directions for my research, and for being able to continue my work on this project.

Many thanks also to Jaap Velthuis, for his supervision and life long lessons at the beginning of my PhD.

Many thanks to Dave Newbold, who for many years has looked out for my best interests, including the time in 'The Robin Hood', where he spoke the words, "Dan, we've got these cubes...".

Special thanks go to Antonin Vacheret, the spokesperson and founder of the SoLid project. Many thanks for the additional supervision you gave at Imperial College, and I look forward to the coming years working together to perform the full experiment!

I am also very grateful for the guidance and help of Nick Ryder throughout my time working on SoLid.

Outside of academia, I am also grateful to many people for their unwavering support, especially: my loving partner Elliott, my mum Siobhan, my dad Alan, my step-mum Vicki, my grand-father Len ('gramps'), and my whole family. Thanks also to Philip Thomas for his support throughout this degree and beyond.

It has also been a pleasure to work within the Bristol University Particle Physics group, for the highly social atmosphere, which I have enjoyed these past eight years. Similarly, I am grateful to the Particle Physics group at Imperial College, who have been extremely welcoming, and I look forward to many more enjoyable times to come.

Author's Declaration

I declare that the work in this dissertation was carried out in accordance with the requirements of the University's Regulations and Code of Practice for Research Degree Programmes and that it has not been submitted for any other academic award. Except where indicated by specific reference in the text, the work is the candidate's own work. Work done in collaboration with, or with the assistance of, others, is indicated as such. Any views expressed in the dissertation are those of the author.

SIGNED: .... DATE:..........................

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The Neutrino sector is one of the least well understood in Particle Physics. The surprise discovery of neutrino oscillations at the beginning of the twenty-first century, leading to the conclusion that neutrinos have mass, has led to a set of new mysteries. The absolute mass of neutrinos is currently unknown; hints of the mass ordering are just becoming available; the existence of Cross-Parity (CP) violation in the neutrino sector is unknown; and the description of the nature of neutrinos (Dirac or Majorana) is still up for debate.

Many experiments have very quickly been able to parametrise the oscillations for the standard three neutrino model [1]. These experiments study oscillations at many different baselines, using different neutrino sources including: solar, atmospheric, beam, radioactive source, and reactor experiments. In doing so, these experiments have also measured a variety of unexpected results. There are tensions between results from several muon neutrino beam experiments, and experiments report significant deficits in the rate of electron anto-neutrinos detected from both radioactive sources and reactors [2]. A fourth sterile neutrino, at a much higher mass than those of the current three neutrinos, can help to resolve these anomalies. The current status and understanding of neutrino oscillations is described in chapter 2.

Several experiments are preparing to search for sterile neutrinos by placing detectors at very short distances from intense compact neutrino sources, such as research reactors [3]. Conducting experiments in these environments can be very difficult due to increased backgrounds from the reactor itself, and cosmic radiation (being near the surface). Detectors must also be more compact than previously, given the limited space available near reactors. The SoLid experiment (**S**hort baseline **O**scillation search using a **Li**thium-6 **d**etector) uses a novel detector technology to tackle these challenges, and this is the topic of chapter 3.

To test the viability of the technology, a large scale prototype (around 20% the planned mass) was constructed and deployed in 2015. It used a configuration of 16×16×9 cubes. The following three chapters are dedicated to results from this prototype. Chapter 4 describes the reconstruction algorithms developed to analyse the prototype data. This begins with analysing raw data (SiPM waveforms), and ends with fully formed IBD candidates. As this is the first attempt at performing reconstruction using this technology at the large scale, this work provides the blueprint for the data reconstruction at the full experiment. These reconstructed events are then used as input to the inverse beta decay (IBD) analysis described in chapter 6. Selection criteria, using multiple techniques, are used to enhance the signal-to-noise ratio of the dataset, and a neutrino signal is searched for using comparisons of reactor on and off data. Critical to this search, and other analyses that use the IBD candidates as input, is accurate energy calibration of the detector. Cosmic muons that can be tracked through the detector can be used as a calibration source, allowing the energy response of individual cubes to be calibrated to a precision of around 2%. This is described in chapter 5.

Finally, this work has informed many design aspects of the full detector, which is currently being constructed. This detector will use a configuration of 16×16×50 cubes. Chapter 7 ends by briefly describing a selection of design decisions that were informed by this research, and also presents the expected sterile search sensitivity. The entire body of work is summarised in the conclusions in chapter 8.

Chapter 2

Neutrino Oscillations

2.1 Introduction

This chapter provides the experimental context and theoretical description of neutrino oscillation, to set the scene for SoLid. The introduction states how neutrinos are described in the Standard Model of Particle Physics. The chapter continues by briefly sketching the discovery of oscillations by the solar and atmospheric neutrino experiments in section 2.2. The theoretical description of neutrino oscillations is described in section 2.3, and up-to-date measurements of the oscillation parameters are presented in section 2.4. Emphasis is given to the current generation of reactor experiments, and the associated anomalies with their measurements. The chapter ends by describing how these anomalies could be better understood by performing measurements of the Uranium-235 neutrino energy spectrum, and also comments on whether sterile neutrinos are a plausible explanation in section 2.5. The current experimental status of sterile searches, as well as future experiments such as the next generation of reactor based experiments, will be presented.

2.1.1 Neutrinos in the Standard Model

The Standard Model (SM) of particle physics describes the interactions of particles via the strong, weak and electromagnetic interactions. It is a quantum field theory based on the gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where C, L, and Y refer to the colour charge, left-handed chirality and weak hyper-charge respectively. The $SU(3)_C$ term describes strong interactions, and the $SU(2)_L \times U(1)_Y$ terms describe the weak and electromagnetic forces. It describes interactions between three kinds of particles:

- **Gauge Bosons**: particles with spin 1. These are the generators of the gauge

Fermions	First Family	Second Family	Third Family
Quarks	$(u,\ d)_L$	$(c,\ s)_L$	$(t,\ b)_L$
	u_R, d_R	c_R, s_R	t_R, b_R
Leptons	$(e^-,\ \nu_e)_L$	$(\mu^-,\ \nu_\mu)_L$	$(\tau^-,\ \nu_\tau)_L$
	e_R	μ_R	τ_R

Table 2.1: The Standard Model Fermions.

symmetries.

- **Fermions**: particles with half-spin. Fermions are sub-divided into two categories: quarks and leptons, as shown in table 2.1. The leptons (as with the quarks) can be further sub-categorised by their chirality:
  - **Left-handed leptons**: described by doublets of a charged lepton and it's corresponding neutrino (e.g. for the electron family: $(e^-,\ \nu_e)_L$). This notation is used to simplify the descriptions of lepton weak interactions in the SM lagrangian [5].
  - **Right-handed leptons**: described by singlets (e.g. e_R^-). Unlike the quark sector, where for each left-handed quark there exists a right-handed counterpart, the Standard Model explicitly does not describe right-handed neutrinos.
- **Higgs Boson**: particle with spin 0. A product of the Higgs mechanism, which is one explanation that gives rise to massive gauge bosons and fermions.

Since neutrinos are neutral leptons, they interact only via the weak force. The weak force itself only interacts with left-handed chiral fields; right-handed neutrinos have never been observed experimentally, and their existence could lead to many solutions to current open questions in Particle Physics [9]. On the other hand, left-handed neutrinos can interact with matter in one of two ways (shown in figure 2.1):

- **Neutral current interactions**: mediated by the exchange of a Z boson. Examples of these interactions include neutrino-electron scattering. The resulting energy change can increase the energy and momentum of the outgoing charged lepton.
- **Charged current interactions**: mediated by the exchange of a W^+ or W^- boson. Since these bosons are charged, and given lepton number is conserved in the SM, the

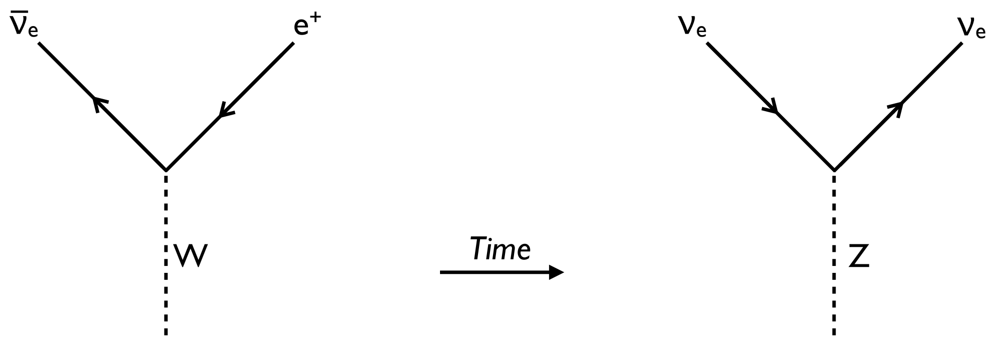

Figure 2.1: Example Feynman diagram vertices of neutrino interactions. Left: Charged current. Right: Neutral current.

neutrino will convert into the corresponding charged lepton. Examples include beta decay.

2.2 Discovery of Oscillations

The discovery of neutrino oscillations by the Sudbury Neutrino Observatory (SNO) [6] and the Super-Kamiokande (Super-K) [7] experiments in the early 2000s was awarded the Nobel Prize for Physics in 2015. First indications that neutrinos oscillate were seen as early as the late 60s. Experiments that measured the solar flux of neutrinos, such as 'The Davis Experiment' (AKA 'Homestake' [8]) found a significant deficit in the rate of solar electron neutrinos of around two thirds. This become known as the 'Solar neutrino problem', which was finally resolved by the unambiguous observation of neutrino oscillations by SNO and Super-K.

2.2.1 Solar Neutrinos

The Sun is a large source of neutrinos,. Those that have large enough energy to be detected by experiments on the Earth ($\gtrsim$ 1MeV) are electron neutrinos, typically produced from the decay of boron-8, as part of the *pp* chain [5]. Many of the older solar neutrino experiments from the twentieth century, including the Homestake experiment, used radiochemical detection techniques, such as electron neutrino capture on chlorine:

$$\nu_e + {}^{37}\mathrm{Cl} \rightarrow {}^{37}\mathrm{Ar} + e \tag{2.1}$$

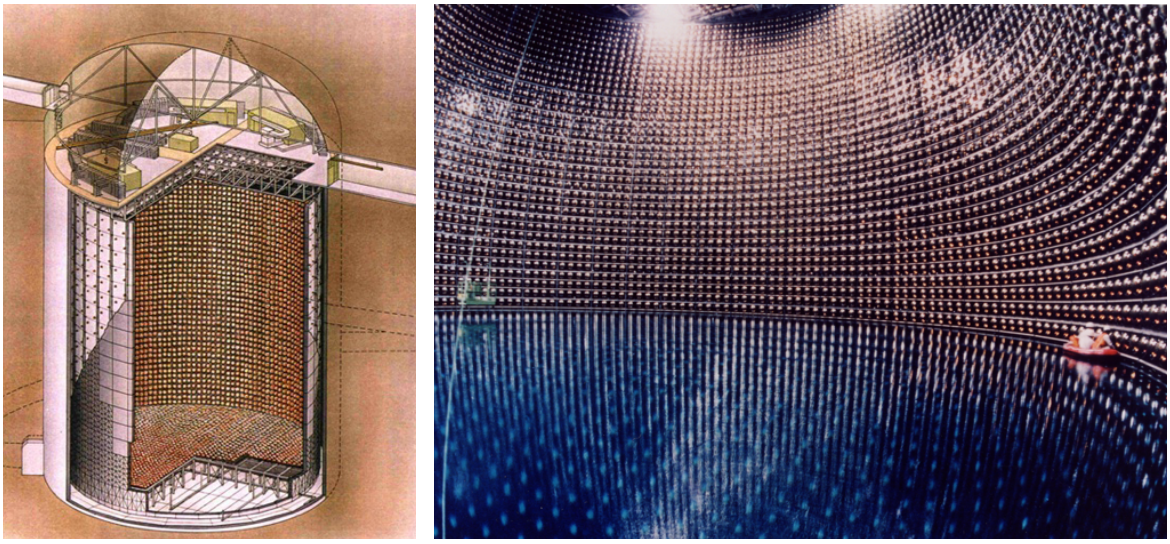

Figure 2.2: The Super Kamiokande (Super-K) experiment, based under Mount Ikeno, 1.0 km underground. Left: Diagram of the detector under the mountain. Right: Picture from inside the detector volume during cleaning.

The subsequent argon atoms can be detected once isolated and upon decaying back into chlorine via electron capture. Some other experiments, such as Kamiokande [10], detected solar neutrinos using Cherenkov radiation. In neutrino reactions that involve a high energy charged lepton in the final state, such as elastic neutrino scattering (neutral current): $\nu + e \rightarrow \nu + e$, the lepton can be detected due to the emission of Cherenkov radiation, when traversing a medium such as water.

The hypothesis that neutrinos oscillated into other flavours was first proposed in the late 1960s by Bruno Pontecorvo [14]. He proposed that electron neutrinos emitted by the Sun may transform into other neutrino flavours at the distance scale (AKA baseline) of the Earth's distance from the Sun, thus reducing the flux of electron neutrinos at the Earth. Since neutrino experiments at the time were mostly insensitive to these other flavours, firm evidence for oscillations came later in the early 2000s, with the Super-K (see figure 2.2) and SNO (see figure 2.3) experiments. Super-K mostly studied atmospheric neutrinos (see below), but was also sensitive to solar neutrinos.

SNO used heavy water as a Cherenkov medium, providing sensitivity to both charged and neutral current neutrino interactions. Since solar neutrinos have low energies compared to the mass of the muon and the tau, charged current interactions could only involve electron neutrinos. An electron neutrino can convert a neutron from the deuteron into a proton

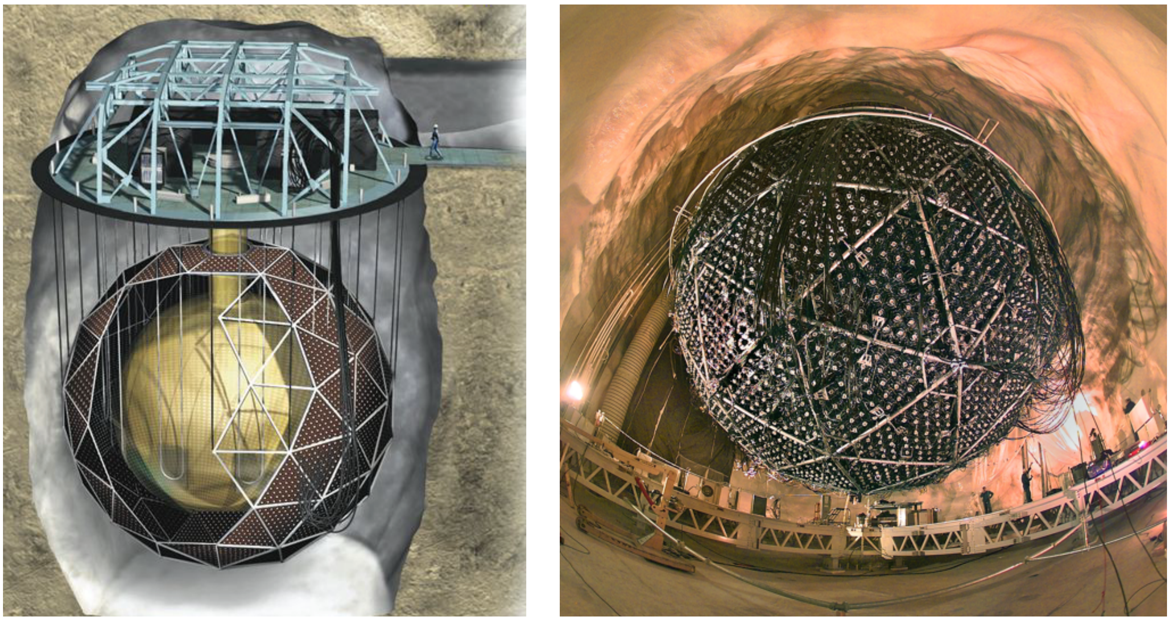

Figure 2.3: The Sudbury Neutrino Observatory (SNO), based in the INCO Creighton mine, 2.1 km underground. Left: Digram of the detector hall in the mine. Right: Picture of the detector volume and surrounding photomultiplier tubes.

via: $\nu_e + d \rightarrow p + p + e$. The emitted electron is then detected via Cherenkov radiation. Both experiments measured the total neutrino flux from neutral current interactions, which matched theoretical predictions of the total solar neutrino rate. Combined with the normalisation provided by the charged current interactions at SNO, this was strong evidence of solar neutrino oscillations.

2.2.2 Atmospheric

Super-K was also able to observe atmospheric neutrino oscillations. Cosmic radiation interactions with the upper atmosphere produce muons and pions that subsequently decay (e.g via $\pi^+ \rightarrow \mu^+ + \nu_\mu$ and $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu_\mu}$), resulting in a large flux of neutrinos produced in the atmosphere. Super-K, like Kamiokande, had the ability to measure the directionality of neutrinos, such as distinguishing an atmospheric neutrino travelling downwards, from a neutrino produced the other side of the Earth travelling upwards. Neutrinos were detected by neutral current intereactions between the neutrino and the Cherenkov medium - the particles in the final state, with increased energy, emmited Cherenkov light that can be detected. Super-K was able to show that the neutrino flux for muon neutrinos depended on zenith angle, in a way consistent with neutrino oscillations at the baseline of the Earth's diameter - a vastly different baseline compared to the solar neutrino problem [13].

2.3 Mixing Parameterisation

The SM states that neutrinos interact with matter as flavour eigenstates: $|\nu_e\rangle, |\nu_\mu\rangle, |\nu_\tau\rangle$. Since neutrinos oscillate during their propagation, as described by the Schrödinger equation, the flavour eigenstates cannot be the eigenstates of the free-particle Hamiltonian [5] (i.e the mass eigenstates). The theoretical extension to the Standard Model for describing neutrino oscillations was formulated by Maki, Nakagawa, Sakata and Pontecorvo. They proposed that the mass eigenstates: $|\nu_1\rangle, |\nu_2\rangle, |\nu_3\rangle$ are a linear superposition of the flavour eigenstates, and the two sets are related by the PMNS matrix U:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \end{pmatrix} \tag{2.2}$$

This is analogous to mixing in the quark sector, where the quark mass eigenstates are related to the quark weak interaction eigenstates by the CKM matrix. Assuming the PMNS matrix is unitary, the matrix can be parametrised by four real parameters: θ_{12}, θ_{23}, θ_{13} and δ, in the form [1]:

$$U = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix} \tag{2.3}$$

where $c_{ij} = \cos\theta_{ij}$ and $s_{ij} = \sin\theta_{ij}$. The current experimental measurements of these four parameters is described in section 2.4, and are consistent with the matrix being unitary.

2.3.1 Two Flavour Case

It is useful to study a simplified case of the above oscillation formalism to demonstrate some of the consequences of neutrion oscillations. Consider the case of just two neutrino flavours $|\nu_e\rangle, |\nu_\mu\rangle$ and two mass eigenstates $|\nu_1\rangle, |\nu_2\rangle$. At a given time, the oscillations can be described by a single angle θ [5]:

$$|\nu_1\rangle = \cos\theta\,|\nu_\mu\rangle - \sin\theta\,|\nu_e\rangle \text{ and } |\nu_2\rangle = \sin\theta\,|\nu_\mu\rangle + \cos\theta\,|\nu_e\rangle \tag{2.4}$$

Equivalently:

$$|\nu_e\rangle = \cos\theta\,|\nu_1\rangle + \sin\theta\,|\nu_2\rangle \;\text{ and }\; |\nu_\mu\rangle = -\sin\theta\,|\nu_1\rangle + \cos\theta\,|\nu_2\rangle \tag{2.5}$$

The Schrödinger equation can be solved to describe the propagation of the mass states as they travel in time t:

$$|\nu_i(t)\rangle = |\nu_i(t=0)\rangle\, e^{-iE_it} \tag{2.6}$$

where E_i is the energy of the ith neutrino (using natural units). Therefore, for the propagation of the flavour states:

$$|\nu_e(t)\rangle = \cos\theta e^{-iE_1t}\,|\nu_1\rangle + \sin\theta e^{-iE_2t}\,|\nu_2\rangle \tag{2.7}$$

$$|\nu_\mu(t)\rangle = -\sin\theta e^{-iE_1t}\,|\nu_1\rangle + \cos\theta e^{-iE_2t}\,|\nu_2\rangle \tag{2.8}$$

Assuming the particle starts as an electron neutrino: $|\nu_1(t=0)\rangle = |\nu_e\rangle$. Solving for the probability that the neutrino has converted into a muon neutrino $P_{\nu_e\to\nu_\mu}$ gives (after re-arranging):

$$P_{\nu_e\to\nu_\mu} = |\,\langle\nu_\mu|\nu_\mu\rangle\,|^2 = \left(\sin(2\theta)\sin\left(\frac{E_2-E_1}{2}t\right)\right)^2 \tag{2.9}$$

$P_{\nu_e\to\nu_\mu}$ is also known as the 'appearance probability' of the muon. For a highly relativistic particle, $E_i \approx p_i + \frac{m_i^2}{2p_i}$ (binomial expansion), where p_i is the magnitude of the momentum of the ith particle, and m_i is it's mass. By also substituting $t = L$ (natural units), where L is the distance travelled by the neutrino since $t = 0$, equation 2.9 can be re-written as [5]:

$$P_{\nu_e\to\nu_\mu} = \left(\sin(2\theta)\sin\left(\frac{m_2^2-m_1^2}{4E}L\right)\right)^2 \tag{2.10}$$

This equation has profound consequences. It can be seen that the following two conditions must be satisfied for neutrinos to oscillate:

- $\sin(2\theta)$ is non-zero.
- There is an energy (i.e. mass) difference between the two mass eigenstates. Therefore, neutrinos are required to have non-zero mass.

It can also be seen that in addition to the angle θ and difference in mass (AKA mass splitting: $\Delta m^2 = m_2^2 - m_1^2$), the wavelength of the oscillation also depends on the ratio of the distance from the neutrino source and the energy of the neutrino: L/E. This has many consequences for the design of oscillation experiments, which are required to tune this ratio for best sensitivity for measurements.

It can be shown that the general appearance (and disappearance) probabilities in the case of three neutrinos take the form:

$$P_{\nu_\alpha \to \nu_\beta} = \sum_{i,j=1}^{3} U_{\alpha i}^* U_{\alpha j}^* U_{\beta i} U_{\beta j} \exp\left(-i\frac{\Delta m_{ij}^2 L}{2E} \right) \tag{2.11}$$

where Greek indices represent the neutrino flavours. There are two independent mass splittings: $\Delta m_{21}^2 = m_2^2 - m_1^2$ and $\Delta m_{32}^2 = m_3^2 - m_2^2$. Oscillation experiments can therefore study the mass differences between the neutrinos. This includes the mass ordering of the neutrinos - there is currently little evidence to distinguish the mass orderings $m_1 < m_2 < m_3$ and $m_3 < m_2 < m_1$ (n.b $m_2 > m_1$ by definition). This is known as the neutrino mass hierarchy problem, also shown in figure 2.4. In addition to experiments such as KATRIN, which aim to find the absolute mass of neutrinos, oscillation experiments can also help solve the hierarchy problem. Experiments studying appearance probabilities, which depend on both Δm_{21} and Δm_{31}, such as T2K (see below) are sensitive to the mass ordering of neutrinos.

2.3.2 Neutrino Mass - Dirac or Majorana?

Dirac Mass Term

The SM does not describe massive neutrinos. The mass of other fermions in the SM, which satisfy the Dirac Equation, are described with terms in the SM Lagrangian $\mathcal{L}_{SM}$ that take the general form: $\mathcal{L}_{Mass}^{Dirac} = -m\psi\bar{\psi}$ [5], where m is the mass of the particle, ψ is the Dirac

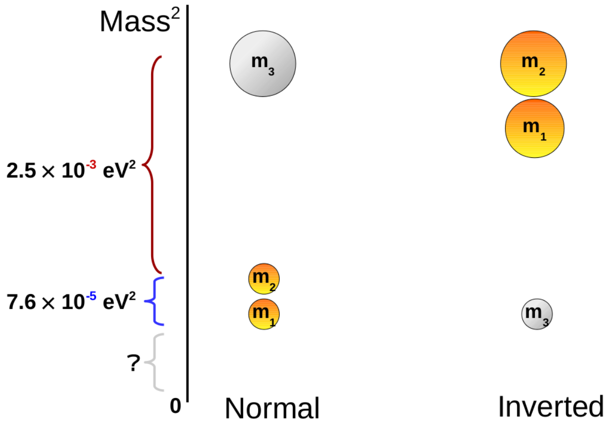

Figure 2.4: The two possible neutrino mass hierarchies. *Credit: Super-K Collaboration.*

spinor, and $\bar{\psi}$ is the adjoint spinor. In order to describe the mass of a neutrino with a Dirac mass term, it is necessary to introduce right-handed neutrinos to the SM, such that the mass term takes the form [11]:

$$\mathcal{L}^{Dirac}_{Mass} = -m(\bar{\nu}_R \nu_L + \bar{\nu_L} \nu_R) \tag{2.12}$$

This model is sometimes referred to as the 'minimally extended Standard Model'. The neutrino right-handed fields are singlets under all SM interactions, and are sometimes referred to as sterile (since they do not interact via any of the SM forces). There is no experimental evidence for the existence of right-handed neutrinos. Lepton flavour conservation is required for neutrino masses to be described by Dirac mass terms [11].

Majorana Mass Term

An alternative approach was suggested by Ettore Majorana in 1937 [11]. He proposed that unlike other SM fermions, the neutrino could be its own anti-particle (i.e invariant under CPT transformation). The Majorana mass term takes the form:

$$\mathcal{L}^{Majorana}_{Mass} = -m(\bar{\nu^C_L} \nu_L + \bar{\nu_L} \nu^C_L) \tag{2.13}$$

where the superscript C denotes the charge-conjugate field. One consequence of Majorana neutrinos is that lepton flavour number is not a conserved quantity, which is beyond the SM.

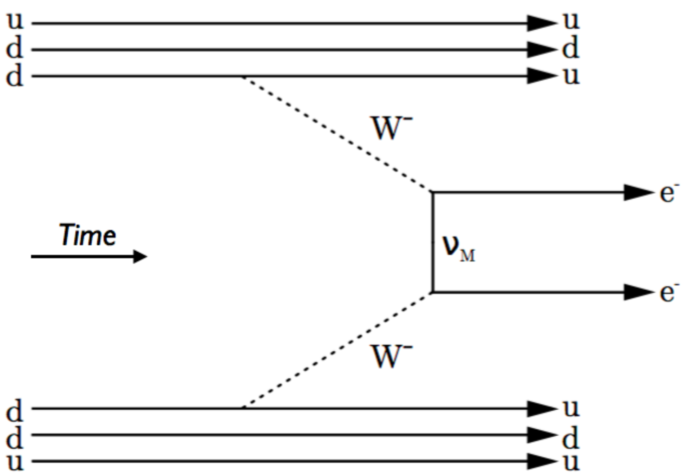

Figure 2.5: Feynam Diagram showing neutrinoless double beta decay. Two neutrons transform into two protons due to the exchange of a Majorana neutrino [12].

This can allow for decays such as neutrinoless double beta decay $0\nu\beta\beta$: $2n \rightarrow 2p + 2e^-$ - see figure 2.5. Experiments use this fact to ascertain the nature of neutrinos, such as the SuperNEMO experiment - see [16].

2.4 Measurements

There have been many experiments to precisely measure the six different oscillation parameters: Δm_{21}, Δm_{32}, θ_{12}, θ_{23}, θ_{13} and δ. They all study the appearance and disappearance probabilities, as a function of distance and energy, from different neutrino sources: solar, atmospheric, beam and reactor neutrinos. The event rate of neutrinos that can be detected from these sources varies significantly, and examples from different experiments are shown in table 2.2. Some of these examples are also used below to show various measurements of the different oscillation parameters. The measurements themselves (from global fits) are summarised in table 2.3.

Appearance and disappearance studies are very complementary: it can be shown that the disappearance probability of a particular neutrino flavour depends on a single row of the PMNS matrix [2], allowing for certain parameters to be measured independently of the other parameters (such as the case with θ_{13} and reactor based experiments); whereas appearance probability depends on many more PMNS matrix elements, allowing experiments to study multiple oscillation parameters at once (such as the beam experiments with both θ_{13} and

θ_{23}).

2.4.1 Solar - Δm^2_{21} and θ_{12}

Solar neutrino experiments and long baseline reactor neutrino experiments are especially sensitive to the parameters Δm^2_{21} and θ_{12} [7]. For the solar neutrino experiments, $L/E \sim 10^{10}$ km/GeV. For the KamLAND experiment, which uses reactor neutrinos from a total of 55 reactors[1] in Japan at various baselines, $L/E \sim 10^4$ km/GeV. KamLAND could only detect electron anti-neutrinos, with an average energy of 3.6 MeV [24], and so studied the survival probability $P_{\bar{\nu_e}\to\bar{\nu_e}}$ [25] (see reference for derivation details):

$$\begin{aligned} P_{\bar{\nu_e}\to\bar{\nu_e}} &= 1 - \cos^4\theta_{13}\sin^2(2\theta_{12})\sin^2\left(\Delta m^2_{21}\frac{L}{4E}\right) \\ &\quad - \sin^2(2\theta_{13})\left(\cos^2\theta_{12}\sin^2\left(\Delta m^2_{31}\frac{L}{4E}\right) + \sin^2\theta_{12}\sin^2\left(\Delta m^2_{32}\frac{L}{4E}\right)\right) \end{aligned} \tag{2.14}$$

Given the small value of θ_{13}, and the large value of L/E for KamLAND and the solar experiments, the last term can be neglected, and $P_{\bar{\nu_e}\to\bar{\nu_e}}$ can be approximated as [25]:

$$P_{\bar{\nu_e}\to\bar{\nu_e}} \approx 1 - \sin^2(2\theta_{12})\sin^2\left(\Delta m^2_{21}\frac{L}{4E}\right) \tag{2.15}$$

KamLAND was able to experimentally verify this relation, as shown in figure 2.6. A combined analysis from KamLAND and the solar neutrino experiments (Super-K and SNO) to measure θ_{12} and Δm_{21} is shown in figure 2.7.

2.4.2 Beam and Atmospheric - θ_{23}, Δm^2_{32} and δ

Experiments using beam or atmospheric neutrinos have an L/E ratio suitable for studying θ_{23} and Δm^2_{32}. Both types of experiment use muon neutrinos. For beam experiments, a muon neutrino beam can be produced by aiming a proton beam at a target. The resulting short lived particles, such as pions and kaons, mostly decay to a (anti) muon and muon (anti) neutrino, producing a neutrino beam. Using T2K as an example, the energy of the neutrinos is peaked at 600 MeV [27], and the baseline is 295 km, giving $L/E \sim 2$ MeV/km. The survival probability of muon neutrinos at this L/E can be approximated as [2]:

[1] Reactor neutrino experiments are described in more detail later in section 2.4.3.

Experiment	Neutrino Source	Active Mass (T)	Depth Beneath Surface (m)	Event rate (/year)
Super-K	Solar and Atmospheric	50 k	1.0 k	8.3 k
SNO	Solar and Atmospheric	1.0 k	2.1 k	4.6 k
NOνA	Beam	300 (near), 14 k (far)	0 (near and far)	$\mathcal{O}(100)$ (far)
T2K	Beam	N/A (near)	0 (near)	$\mathcal{O}(100)$ (far)
KamLAND	Reactor (long baseline)	1.0 k	1.0 k	180
Daya-Bay	Reactor (short baseline)	(8×) 20	250 - 860	700 k (total)
RENO	Reactor (short baseline)	(2×) 17	70 (near), 200 (far)	230 k (near), 24 k (far)
Double-Chooz	Reactor (short baseline)	(2×) 8.0	60 (near), 150 (far)	100 k (near), 15 k (far)
SoLid	*Reactor (very short baseline)*	*2.0*	*0.0*	*100 k*
SOX	*Source (very short baseline)*	*280*	*1.2 k*	*10 k*

Table 2.2: Comparison of different neutrino oscillation experiments using different sources and baselines. All values to two significant figures. The event rates are taken as that used for official physics analysis. The T2K near detector is composed of multiple sub-detectors using different detector technologies, and assigning an active mass is not a fair comparison. The far detector of T2K is Super-K. SoX uses the Borexino detector. Future experiments are shown in *italic*, and these values are expectations. Sources: [6, 7, 15, 17, 18, 19, 21, 22].

Parameter	Measurement
Δm^2_{21}	$(7.3 \pm 2.5) \times 10^{-5}$ eV^2
Δm^2_{32}	$(2.50 \pm 0.83) \times 10^{-3}$ eV^2
$\sin^2 \theta^2_{12}$	0.30 ± 0.10
$\sin^2 \theta^2_{23}$	0.44 ± 0.17
$\sin^2 \theta^2_{13}$	0.0214 ± 0.0072
δ (π)	1.34 ± 0.73

Table 2.3: Global fit of current measurements of the neutrino oscillation parameters - see [1]. Experiments performing these measurements are described in section 2.4. Assumes normal mass hierarchy. Errors are 1σ.

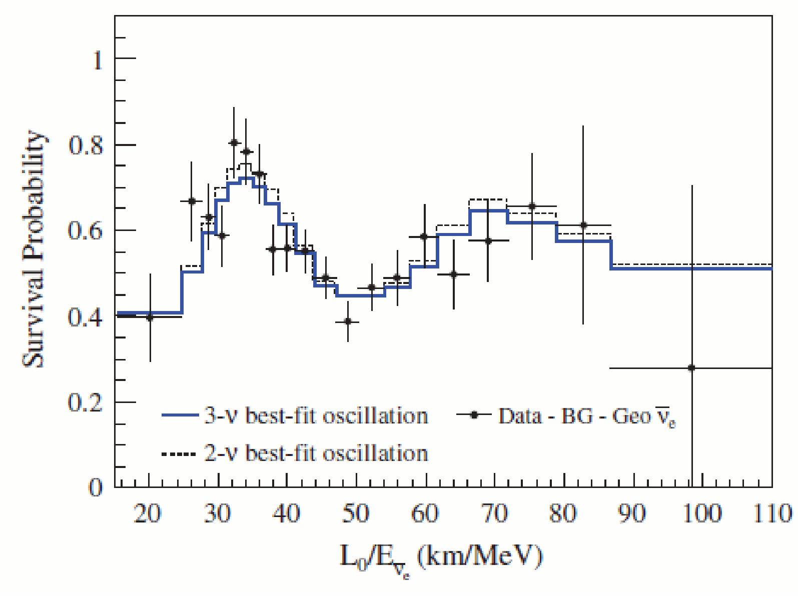

Figure 2.6: Oscillations in L/E measured by the KamLAND experiment [26].

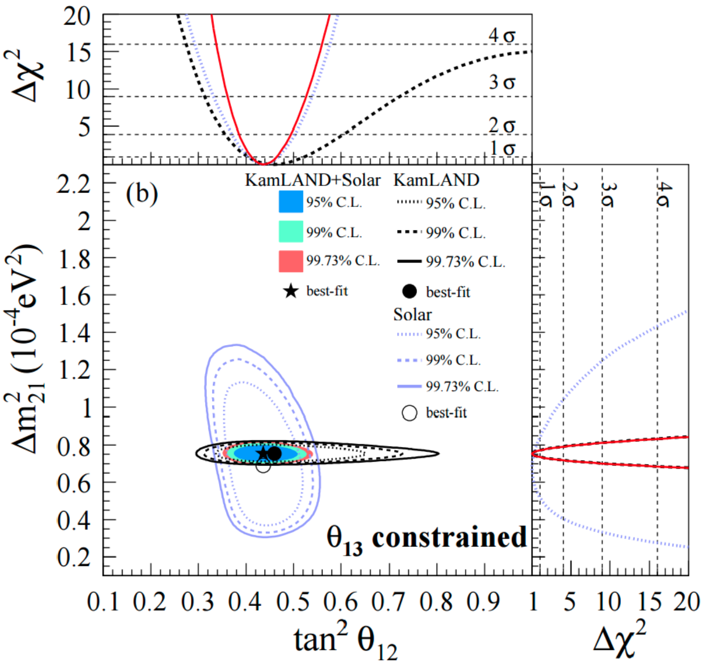

Figure 2.7: Combined analysis of KamLAND and the solar neutrino experiments to find Δm^2_{21} and θ_{12} [26].

$$P_{\nu_\mu \to \nu_\mu} \approx 1 - \left(\cos^4\theta_{13}\sin^2(2\theta_{23}) + \sin^2(2\theta_{13})\sin^2\theta_{23}\right)\sin^2\left(\Delta m^2_{32}\frac{L}{4E}\right) \tag{2.16}$$

Beam experiments can therefore perform precision measurements of θ_{23} and Δm^2_{32}. Some of the most recent results, from multiple experiments, are shown in figure 2.8.

By studying electron neutrino appearance, beam experiments are also able to measure other oscillation parameters, such as θ_{13} and δ [28]. δ is also referred to as the CP violating term - it can be shown that a non-zero value of δ is required to observe CP violation in neutrinos [29]. The appearance probability of an electron neutrino given a muon neutrino in the initial state $P_{\nu_\mu \to \nu_e}$ can be shown to depend on all of θ_{13}, θ_{23}, Δm^2_{31}, δ, θ_{12} and Δm^2_{21}. One consequence of this includes $P_{\nu_\mu \to \nu_e} \neq P_{\bar{\nu_\mu} \to \bar{\nu_e}}$ if and only if $\sin\delta \neq 0$. An example measurement of δ from T2K is shown in figure 2.9. A small preference for non-zero δ is observed for both neutrino mass hierarchies, although this is not yet statistically significant beyond the 90% confidence level.

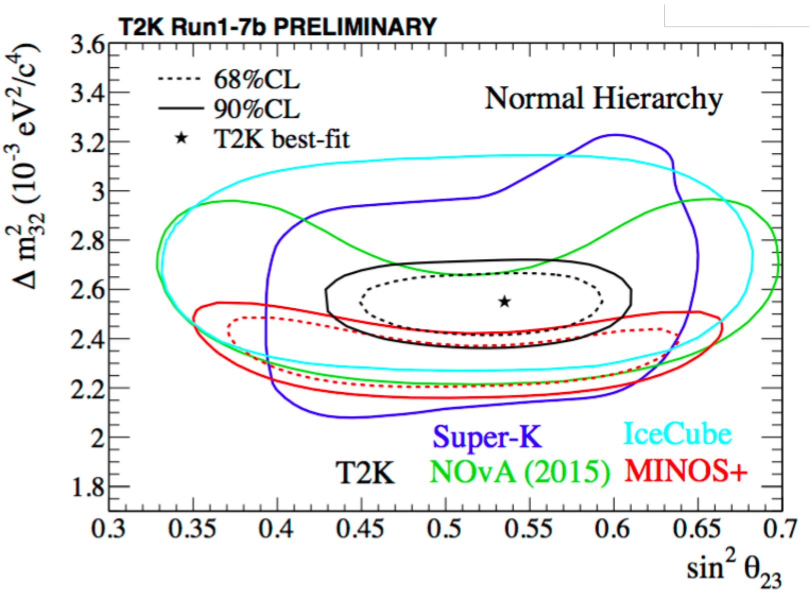

Figure 2.8: Recent measurements of θ_{23} and Δm^2_{32} from atmospheric and beam experiments (produced by T2K [28]).

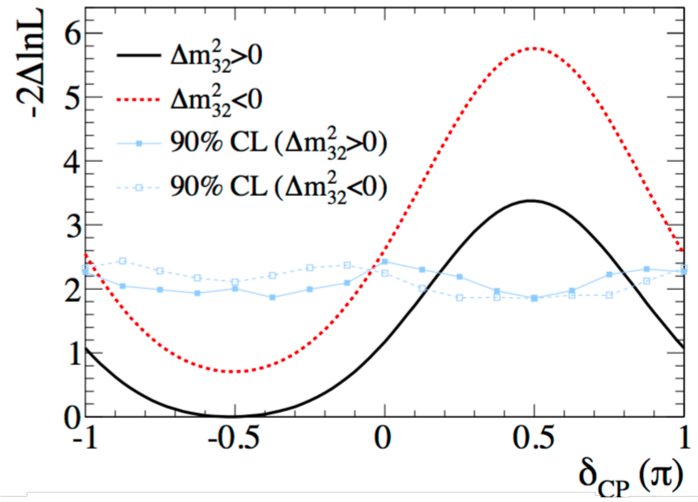

Figure 2.9: Measurement of δ by the T2K experiment, found by studying electron neutrino appearance from a muon neutrino beam. The value of θ_{13} measured by the reactor experiments is used as input [1].

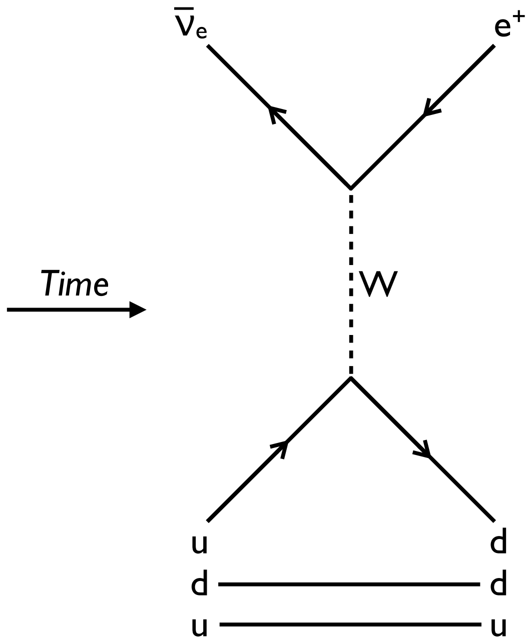

Figure 2.10: Feynman diagram of inverse beta decay.

2.4.3 Reactor - θ_{13}

Both atmospheric and solar neutrino experiments are typically built far underground, to reduce the large backgrounds from cosmic radiation. This is required due to the small flux of neutrino interactions at these experiments. Nuclear reactors are another intense source of neutrinos, and the neutrino flux can be many orders of magnitude higher for experiments using a reactor as a source. This allows experiments to be made more compact, and loosens some of the requirements for shielding against backgrounds.

Inverse Beta Decays

The majority of reactor based (AKA short baseline) neutrino experiments use the same decay channel - inverse beta decays (IBD):

$$\bar{\nu_e} + p \rightarrow n + e^+ \ (E_{\bar{\nu_e}} > 1.805 \text{ MeV}) \tag{2.17}$$

The Feynman diagram is shown in figure 2.10. An electron anti-neutrino is captured by a proton from the detector volume. The outputs of the reaction are a positron, whose energy is highly correlated with the neutrino energy [4], and a neutron. The neutron energy distribution is shown in figure 2.11, and the average is around 20 keV. The positron annihilates with an

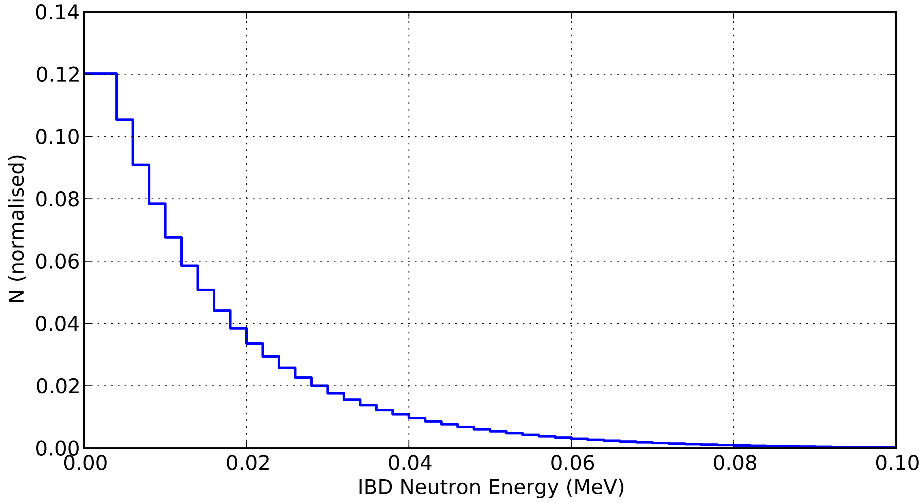

Figure 2.11: Neutron IBD energy distribution from simulations of the SoLid experiment.

electron in the detector, producing two annihilation photons. Signals from the positron and annihilation photons are detected within in a few ns (depending on the scintillator type), and are collectively known as the 'prompt' event. A few microseconds later, after it has thermalisd, the neutron can also be detected - the 'delayed' event. Using these distinctive features, IBD interactions can be well separated from backgrounds. On the other hand, since a co-incidence between two events (prompt and delayed) is required to identify a neutrino, one component of background is accidental candidates - random mismatches of uncorrelated positron and neutron candidates, which can be large in the environment near nuclear reactors.

The Cowan and Reines experiment was the first neutrino experiment to use a nuclear reactor as a source of neutrinos in the 1950's [32]. Based at the Savannah River Plant in South Carolina, the experiment confirmed the existence of the anti-neutrino. The detector used two water tanks as proton targets, which were sandwiched between tanks of cadmium-doped liquid scintillator. The liquid scintillator was sensitive to both neutrons, and photons (via Compton scattering). The neutrons are detected by capture on Cadmiun via the reaction:

$$n + {}^{113}_{48}\mathrm{Cd} \rightarrow {}^{114*}_{48}\mathrm{Cd} \rightarrow {}^{114}_{48}\mathrm{Cd} + \gamma \tag{2.18}$$

The γ can be detected, and it's energy can be used as a tag of the decay.

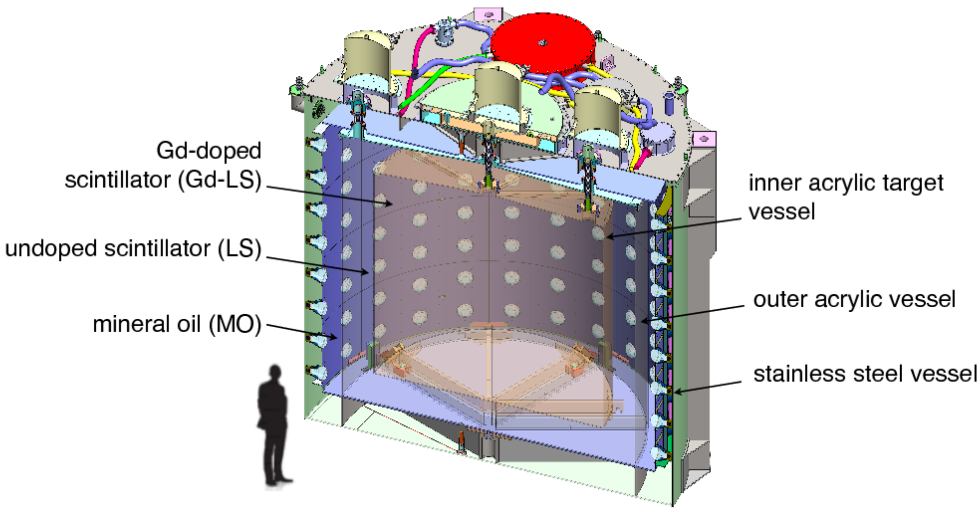

Figure 2.12: A diagram of a module of the Daya Bay experiment.

Current Generation

The current generation of reactor neutrino experiments generally use a similar technology and detector design. Experiments are composed of multiple detector modules, placed at a variety of distances from one or more reactors. An example module from the Daya Bay experiment is shown in figure 2.12. The inside detector volume is Gd-doped liquid scintillator, and provides sensitivity to both EM and neutron interactions. An additional outer layer of non-doped liquid scintillator is also used to detect muons, which acts as a muon veto, and a gamma tagger for any annihilation photons that travel outside of the inner volume.

The outside of the cylindrical vessel is surrounded by photomultiplier tubes to detect the scintillation light, providing a measure of the energy of the position and annihilation photons. Annihilation photons interact with the scintillator via Compton scattering, although in some cases, may exit the detector prior to depositing all its energy. In order to accurately measure the energy of the neutrino, it is critical for these detectors to be able to efficiently detect both the positron and the annihilation gammas (the spatial information is not sufficient to resolve the positron and gammas separately). Similarly, the modules are required to have uniform energy response throughout the detector volume, since the interactions inside the detector cannot be accurately localised.

The experiments themselves:

- **Daya Bay [19]**: consists of eight 20t detector modules, placed at distances ranging from 360m to 1.9km from three pairs of reactors, at the Daya Bay reactor complex in China.

- **Reno [21]**: uses two 16.5t detectors, placed at 290m and 2.4km from a cluster of the six reactors, at the Hanbit nuclear power plant in South Korea.
- **Double-Chooz [22]**: consists of two 8t detectors placed at 400m and 1050m from two reactors, at the Chooz nuclear power plant in France.

All three experiments are able to precisely measure the oscillation angle θ_{13} from electron anti-neutrino disappearance. The L/E for these experiments is ~ 0.5 km/MeV. By revisiting the disappearance probability from equation 2.14, at this scale of L/E, the oscillation is now driven by the θ_{13} term [25]:

$$P_{\bar{\nu_e}\to\bar{\nu_e}} = 1 - \sin^2(2\theta_{13})\left(\cos^2\theta_{12}\sin^2\left(\Delta m^2_{31}\frac{L}{4E}\right) + \sin^2\theta_{12}\sin^2\left(\Delta m^2_{32}\frac{L}{4E}\right)\right) \tag{2.19}$$

It is convenient to define an effective mass splitting m_{ee}:

$$\Delta m^2_{ee} \equiv \cos^2\theta_{12}\Delta m^2_{31} + \sin^2\theta_{12}\Delta m^2_{32} \tag{2.20}$$

which allows equation 2.19 to be written in the same form as other disappearance probabilities presented (to high accuracy), including that used at KamLAND in equation 2.15, and in the two flavour approximation in equation 2.10 [25]:

$$P_{\bar{\nu_e}\to\bar{\nu_e}} \approx 1 - \sin^2(2\theta_{13})\sin^2\left(\Delta m^2_{ee}\frac{L}{4E}\right) \tag{2.21}$$

Like many beam experiments, these experiments also use a configuration of two or more detectors at various distances. This allows ratios of the neutrino flux to be used for physics analysis instead of absolute measurements, removing any dependency from predictions of the absolute neutrino flux and nuclear fission cross sections. This reduces systematic errors (such

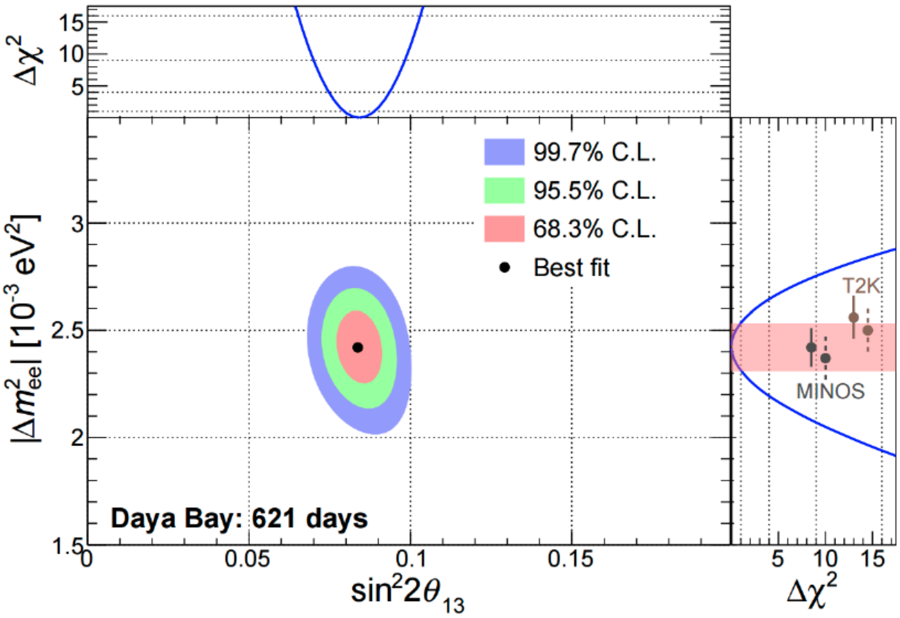

Figure 2.13: Measurement of θ_{13} by the Daya-Bay experiment [20]. Δm_{ee} is an effective mass splitting, defined in the text.

as those due to the reactor anomaly - see below). The θ_{13} result from Daya-Bay is shown in figure 2.13, and the measurement of L/E is shown in figure 2.14.

These experiments, together with past reactor experiments, have measured two main anomalous results: the reactor anomaly, and the 5 MeV distortion.

Reactor Anomaly

In 2012, in prepration for the current generation of reactor experiments, the predicted reactor neutrino flux was re-calculated with increased precision by Muella et al [33], accounting for many more nuclear fission reactions than previously. The new predicted neutrino rate increased by 3.5% compared to the previous calculation. When combined with previous measurements of small deficits in the rate of reactor neutrinos at close distances from reactors, the average deficit is 5.6%, as shown in figure 2.15. This corresponds to a significance at the level of 98.6 % C.L, or 2.7σ [34]. One possible explanation is the existence of a fourth sterile neutrino (see section 2.5). Alternatively, it may be due to unaccounted errors in the calculations of the nuclear fission cross sections, as also demonstrated by the unexpected distortion of the neutrino energy spectrum around 5 MeV.

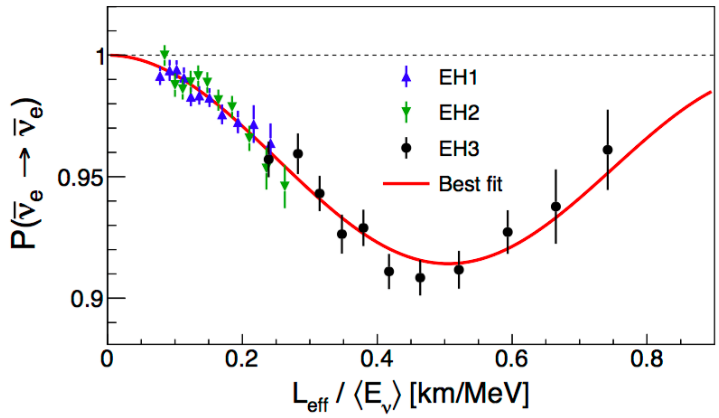

Figure 2.14: Measurement of oscillations in L/E by the Daya-Bay experiment [20]. EH1, EH2 and EH3 are labels of the different sites of detector modules.

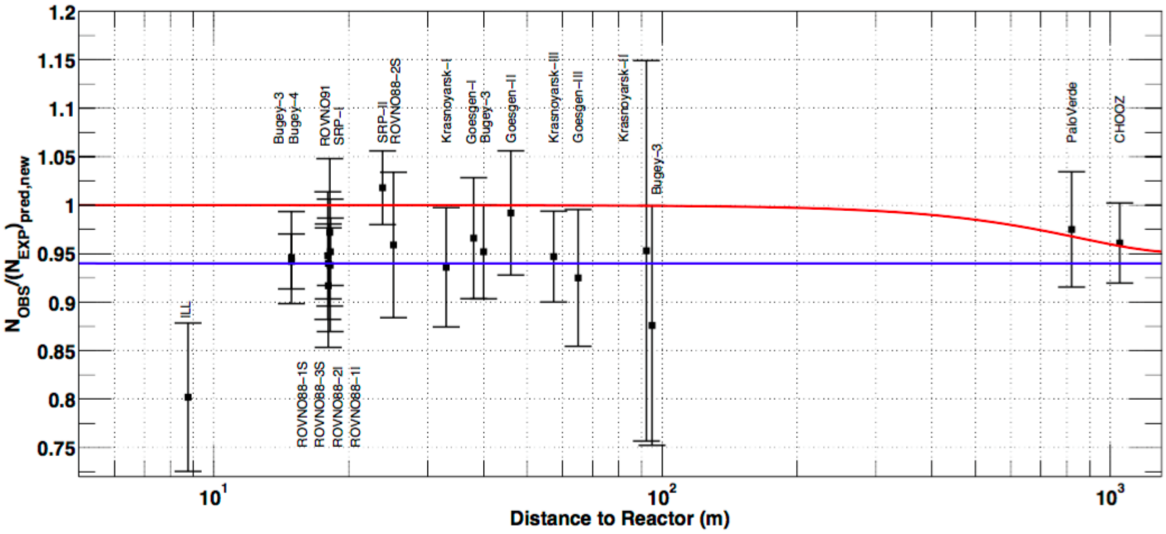

Figure 2.15: Reactor neutrino rate for previous short and very short baseline experiments, using the re-calculated flux prediction from 2012. The red line corresponds to the 3 neutrino hypothesis, and the blue line corresponds to a solution including one sterile neutrino [34].

5 MeV Distortion

All three of the current generation of reactor experiments have observed an unexpected distortion in the shape of the electron anti-neutrino energy spectrum, when compared to theoretical predictions. A measurement by the RENO experiment is shown in figure 2.16. RENO has also shown that the magnitude of the distortion correlates with reactor power [35]. Many possible causes have been suggested, and most assign the distortion to an additional contribution of a single unknown fissile isotope [36]. Due to the similarity of the fuel contents of the three experiments named above, it is difficult assess which isotope is responsible.

Comparisons with future short baseline experiments, which will study reactors with significantly different fuel contents can reveal the source of the distortion. One example comparison is outlined by Huber [36], which compares new results from the NEOS experiment with Daya-Bay. Both experiments use reactors powered by the fission of the isotopes Uranium-235, Uranium-238, Plutonium-239 and Plutonium-241. The respective fission fractions for NEOS are: 0.655, 0.072, 0.235, 0.038; and for Daya-Bay, they are: 0.561, 0.076, 0.307 and 0.056. Comparisons can show that the plutonium-239 and plutonium-241 are unlikely to be the source of the distortion. Huber also points out that future measurements at research reactors (including the Belgian Research Reactor Two), which use almost pure Uranium-235 fission, will help resolve the source of the distortion further. The distortion cannot be explained by a sterile neutrino [34].

2.5 Sterile Neutrinos

One plausible explanation of the Reactor Anomaly is the existence of a fourth neutrino, with additional neutrino oscillations at a reduced value of L/E, requiring a new mass splitting at the scale $\Delta m^2 \sim 1$ eV^2. Experiments at the LEP accelerator measured the number of neutrino flavours that can interact via the weak interaction N_ν by studying Z production in e$^+$ and e$^-$ collisions. The combined measurement gives $N_\nu = 2.984 \pm 0.008$. Therefore, any additional neutrinos cannot interact via weak interactions, and as such are known as ‘sterile’.

In addition to the reactor anomaly, there are also hints from other experiments using a variety of sources and at different L/E scales:

- Radioactive sources placed at solar neutrino experiments, such as SAGE and GALLEX,

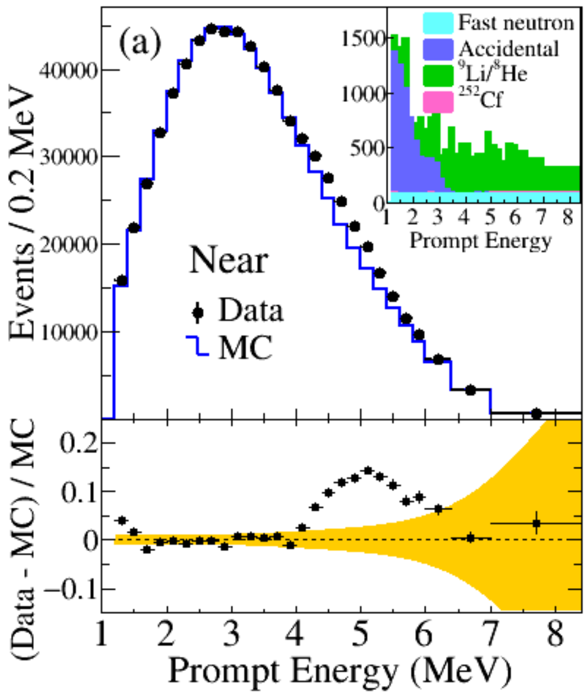

Figure 2.16: Reactor electron anti-neutrino energy spectrum, as measured by the RENO experiment [35]. The source of the distortion in the range between 4-6 MeV is unknown.

also observe a deficit in the rate of neutrino interactions - the 'Gallium Anomaly'. This is shown in figure 2.17, and the significance of this deficit is around 3σ [37].

- The LSND experiment reports observations of $\bar{\nu_\mu} \to \bar{\nu_e}$ oscillations at $E/L \sim 1$ eV^2, which is inconsistent with a 3-neutrino hypothesis [2]. The MiniBooNE experiment has also searched for these transitions, and has observed an unexplained excess at low electron neutrino energies. The results of both experiments are in tension, although can be made consistent in certain sterile neutrino models [2].

There may also be limits on the number of neutrinos from Cosmology, since sterile neutrinos could contribute to the number of degrees of freedom in the early universe, although the extent to which Cosmology favours or disfavours the existence of a sterile neutrino is still under discussion [2].

2.5.1 Theoretical Extension

The oscillation formalism outlined in section 2.3 can be extended to incorporate sterile neutrinos by introducing new mass eigenstates. In the case of one additional sterile neutrino

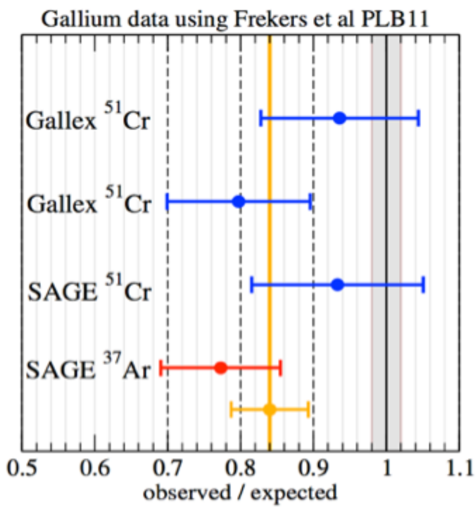

Figure 2.17: Measurements of the neutrino flux for Gallium radioactive source experiments [37].

(AKA the '3+1 model'), the oscillation matrix includes a new row and column for the additional mass state $|\nu_4\rangle$ and sterile flavour state $|\nu_s\rangle$:

$$\begin{pmatrix} \nu_e \\ \nu_\mu \\ \nu_\tau \\ \nu_s \end{pmatrix} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} & U_{e4} \\ U_{\mu 1} & U_{\mu 2} & U_{\mu 3} & U_{\mu 4} \\ U_{\tau 1} & U_{\tau 2} & U_{\tau 3} & U_{\tau 4} \\ U_{s1} & U_{s2} & U_{s3} & U_{s4} \end{pmatrix} \begin{pmatrix} \nu_1 \\ \nu_2 \\ \nu_3 \\ \nu_4 \end{pmatrix} \tag{2.22}$$

In the very short baseline limit, $L/E \sim 10^{-3}$ km/MeV (two orders of magnitude smaller than that of Daya-Bay), the oscillations can be driven by the sterile terms only, giving the familiar form of the electron anti-neutrino disappearance:

$$P_{\bar{\nu_e} \to \bar{\nu_e}} \approx 1 - \sin^2(2\theta_{ee}) \sin^2\left(\Delta m^2_{41} \frac{L}{4E}\right) \tag{2.23}$$

where Δm_{41} is an additional mass splitting, and θ_{ee} is an effective oscillation angle. To demonstrate the effect of these new parameters, an example model is shown in figure 2.18. This shows the expected rate of electron anti-neutrinos as a function of distance from a source, from the very short baseline through to the baseline of the solar experiments. By eye, the data clearly prefer the sterile neutrino hypothesis, although this is expected even were there

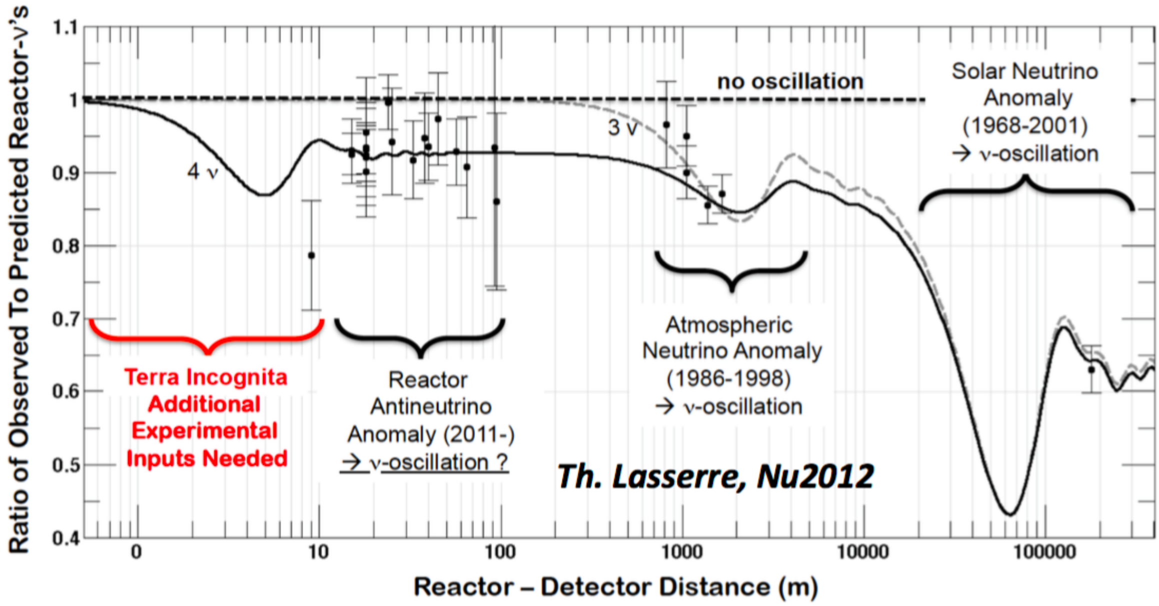

Figure 2.18: Reactor neutrino survival probability as a function of distance from the reactor. The data points represent previous experiments. The dashed gray line shows the expectation for a 3 neutrino hypothesis. The solid black line shows the expectation for an example 3+1 sterile hypothesis.

no sterile neutrinos, since this extension introduces two new free parameters. A fit using the rate measurements from short baseline experiments has been performed by Kopp et. al., who find a best solution:

$$\Delta m_{41} = 0.44\text{eV}^2 \text{ and } \sin^2\theta_{ee} = 0.13 \tag{2.24}$$

The allowed regions (to 90%, 95% and 99% confidence level) are shown in the left of figure 2.19. The best fit solution is shown in the right of figure 2.19, which shows the reactor neutrino rate as a function of distance from the reactor core. Kopp et. al. have also performed fits using the shape of the neutrino energy spectrum (measured by Bugey-3), and this best fit solution is also shown.

2.5.2 Current Search Status

Other experiments can help constrain the sterile neutrino model parameters. For example, studies of muon neutrino disappearance due to oscillations with a sterile mass state can put limits on the sterile neutrino parameters - an example is shown in figure 2.20. This shows the

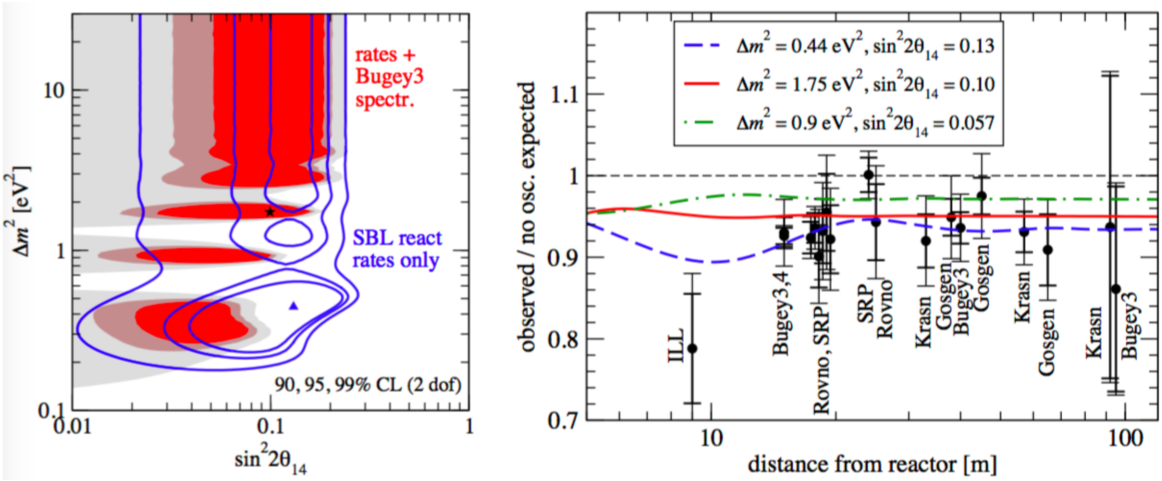

Figure 2.19: Left: Allowed regions of of '3+1' sterile neutrio hypothesis, from a 2D scan over the additional oscillation parameter phase space [2]. Blue shows the fit to the rate measurements reported by short baseline experiments, whereas red also incorporates the neutrino energy spectrum measurement by Bugey-3. Right: Reactor neutrino survival probability, as a function of distance from the reactor, for very short baseline experiments. The best fit using rate information only is shown in **blue**, and the fit also using the Bugey-3 energy spectrum is shown in red **red**.

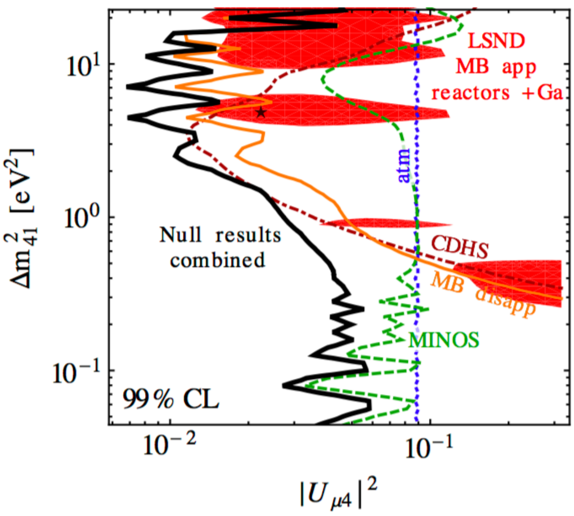

Figure 2.20: Constraints (99% CL) on 3+1 sterile neutrino hypothesis from muon neutrino disappearance experiments (combined in black) [2]. The allowed regions from fits using the reactor neutrino disappearance are shown in red, and are in tension with the muon data.

allowed regions from the reactor experiments, which can be seen to be in tension with the limits from the beam experiments such as MINOS and MiniBooNE. Kopp et. al. conclude from this result that an additional single sterile neutrino cannot explain all anomalies found by oscillation experiments. Similar constraints also arise from muon neutrino appearance studies - see [2] for further details.

2.5.3 Next Generation Reactor Experiments

The next generation of reactor experiments aim to resolve the reactor anomaly and confirm (or refute) the sterile neutrino hypothesis as an explanation. They will search baselines $\mathcal{O}(10)$m, and are referred to as 'very short' baseline neutrino experiments. These experiments will also measure the Uranium-235 neutrino energy spectrum, which will also help to resolve the reported distortion around 5 MeV. There are multiple experiments planned, studying reactors that use a variety of fuel contents. Compared to previous reactor experiments, those at very short baselines face additional and more extreme challenges, including:

- **Compactness**: there are often many obstacles in the environment near reactors, such as building infrastructure for reactor maintenance. The future experiments have been designed to be significantly more compact than those of the current generation (e.g. the active mass of SoLid is over an order of magnitude smaller than that of Double-Chooz. The event rate is compensated by being nearer the reactor).

- **Detector resolution**: as with the current short baseline experiments, good energy resolution is critical to perform an oscillation search. Additionally, position resolution is also crucial. The size of the current short baseline detectors is very small relative to the baseline, and so they use multiple detectors to perform oscillation measurements. This is not the case with the very short baseline experiments, where the size of the detector is the same scale as the baseline. Thus, detectors must also have accurate position resolution to be able to localise the position of the neutrino interactions.

- **Reduced overburden**: as shown in table 2.2, neutrino experiments are often placed underground, to reduce backgrounds that arise from cosmic radiation. Since very short baseline experiments have to be placed very near a reactor, the level of overburden depends on the reactor site, and is often zero.

- **Reactor backgrounds**: the reactor itself can be a source of backgrounds. Reactors output a large flux of gamma rays, which can significantly increase the accidental background.

- **Security**: reactor sites have many security protocols. This includes discouraging the use of flammable liquids, such as some liquid scintillators.

- **Reactor choice**: the reactor itself can influence the experiment. The event rate is proportional to the reactor power. The size of the reactor core should not be large relative to the baseline (as is the case for power reactors), since this will add a large uncertainty to the decay distance measurement. It is also useful for the reactor to have regular on/off cycles, since reactor off datasets are useful for studying some background components.

The next chapter outlines how SoLid has approached these challenges. The competing experiments to SoLid, such as PROSPECT (see figure 2.21) and STEREO (see figure 2.22), have adopted different designs [38]. In summary, the design of future reactor experiments use a different balance of the following parameters:

- Passive shielding.

- Energy and position resolution.

- Scale of segmentation.

- Choice of scintillator.

The future experiments have chosen a different balance - for example:

- STEREO is composed of five large cells using liquid scintillator, providing excellent energy resolution, sufficient position resolution and the ability to remove some backgrounds based on the shape of scintillation signals. STEREO also uses a large amount of passive and active shielding.

- PROSPECT also uses liquid scintillator, placed in an arrangement of 120 bars, offering better segmentation than STEREO. PROSPECT uses reduced passive shielding compared to STEREO, and no active shielding.

- SoLid has a similar amount of shielding to PROSPECT, and uses an arrangement of 13k stacked composite scintillator cubes. This design offers excellent spatial resolution and

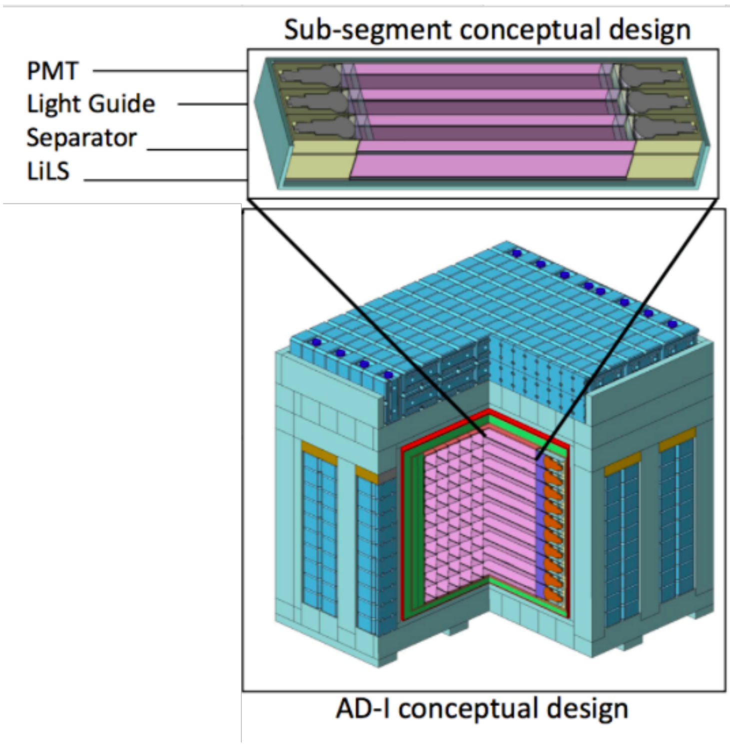

Figure 2.21: The PROSPECT detector design - a very short baseline reactor experiment, using the High Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory (ORNL).

allows for topological measures of events. The ability to remove certain backgrounds based on scintillation signal shape is reduced.

All three experiments expect an IBD efficiency of $\sim 30\%$, a signal to noise ratio $\mathcal{O}(1)$, and plan to begin data taking in 2017. They aim to give an initial sterile neutrino search result within a couple of years. They will all be placed at research reactors. The sensitivity of SoLid for discovering sterile neutrinos is presented in chapter 7.

Source Experiments

Experiments using intense radioactive neutrino sources can also search for oscillations at a very short baseline. An example includes the SOX experiment, which will place an intense $\bar{\nu}_e$ source ^{144}Ce beneath the BOREXINO solar neutrino experiment. Unlike the reactor experiments, SOX will have a minimal cosmic background due to the large overburden. Instead, the challenge it faces is a large gamma-ray background from the source itself. The event rate is also an order of magnitude lower than that expected for the future reactor experiments.

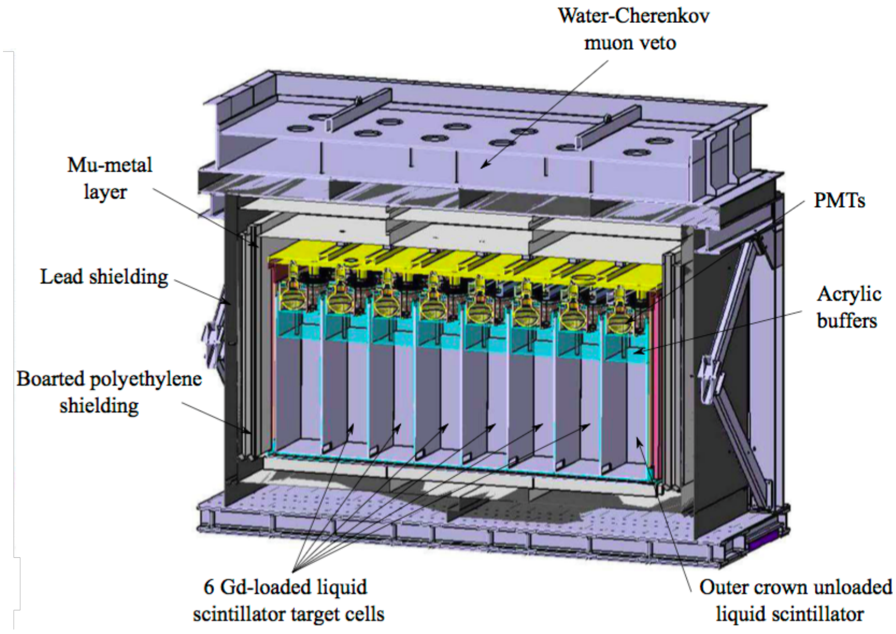

Figure 2.22: The STEREO detector design - a very short baseline reactor experiment using the Institut Laue-Langevin (ILL) reactor in France [39].

Non-proliferation

The compact design of these detectors (SoLid in-particular) means the SoLid detector can also be used as means to promote nuclear non-proliferation, since it is highly desirable to have a system that can remotely verify the use of nuclear materials (such as at a reactor site). In practice, the low mass of the detector results in long exposure times required in order to perform measurements of potential nuclear activity (depending on the distance between the source and the detector), which are sometimes impractical - similar technologies are being explored, see [30].

2.6 Conclusions

The discovery of neutrino oscillations has led to a surge of new experiments in the past two decades. Oscillations are firm evidence that neutrinos are massive, which is not described by the Standard Model of Particle Physics. Experiments have very quickly measured many of the oscillation parameters, using a variety of neutrino sources, and at a range of different baselines. There are still many open questions in the neutrino sector - the absolute mass of neutrinos and their ordering is unknown; we do not understand the mechanism by which neutrinos gain mass in the Standard Model; is there CP violation in the neutrino sector?

Many experiments are under-way to help answer these questions.

Experiments have found a set of anomalous results, such as the reactor and gallium anomalies, which are yet to be resolved. There are also tensions between results from different beam experiments. Measurements of the Uranium-235 neutrino energy spectrum can shed light on the source of these anomalies. One possible explanation is the existence of a sterile neutrino. Many new experiments are beginning to search for neutrino oscillations at the very short baseline, using both reactors and radioactive sources. In addition to searching for sterile neutrinos, the reactor experiments can provide insight to the source of the 5 MeV distortion in the electron anti-neutrino energy spectrum by directly measuring the contribution from Uranium-235. These experiments use different approaches to manage the challenges of performing experiments very near reactors. SoLid, in particular, uses a new technology, and this is the topic of the next chapter.

Chapter 3

SoLid Detector Technology

3.1 Introduction

The SoLid collaboration have developed a novel detector technology to perform a neutrino oscillation search in the challenging environment near nuclear reactors. The technology combines two types of scintillator, allowing for the detection of both neutrons and charged particles, which are sandwiched into a small light-tight cube shaped voxel. These voxels can be stacked to form a compact detection volume, which is suitable for environments close to nuclear reactors. The light tightness allows light from particle interactions to be localised with high precision, and the topological information provided by the high segmentation provides new handles for discriminating signal and background. This chapter explains the technology in more detail, and the principles for it's use to detect anti-neutrinos. This also includes detailed descriptions of the most recent large scale prototype, and the design of the first phase of the experiment, which is currently under construction. Finally, the reactor itself is described.

3.2 Detection Principle

The entire SoLid detector is formed of thousands of small cubes. Each has a side length of 5 cm, as shown in the top diagram of figure 3.1, and in the picture in figure 3.2. The bulk is Polyvynil-Toluene (PVT) - a plastic scintillator that is sensitive to ionising particles, such as positrons or cosmic muons, and has both high light yield and optical transparency [4]. Each cube is polished, then wrapped in a layer of tyvek to increase the internal reflection of the cube, and prevent light leakage to other nearby cubes. Square grooves are cut in the sides of the cube to allow the placing of wavelength shifting optical fibres for light extraction.

Each cube is coupled to four fibres, with two fibres per vertical and horizontal direction. This allows signals to be localised to specific cubes (in 3D). The scintillation photons reach the optical fibres with a wavelength in the blue visible spectrum, around 420 nm. To enter the optical fibres, the angle of incidence is required to be close to perpendicular with the surface of the fibre. The optical fibres are wavelength shifting, and the incident light is re-emitted as green optical photons. Wavelength shifters are used since the loss of light due to attenuation is smaller for green light than for blue. Also, since the shifting process randomises the directionality of the incoming photons, the angular bias towards the normal is removed, which increases the probability of internal reflection, and thus the probability of light reaching the ends of the fibre before absorption.

The surfaces of the fibres are double cladded with an acrylic [31], which has a significantly higher refractive index than the bulk of the fibre. Thus, light travelling inside the fibres is trapped due to total internal reflection, preventing leakage into other cubes along the fibres. Each fibre is read out at one end with a Silicon Photomultiplier (see section 3.2.2). The other end of the fibre is coated with a reflective material.

Two sides of the cube are lined with a sheet of ^{6}LiF:ZnS(Ag), which are 250 μm thick, and provide a means to detect neutrons. The neutron is captured on the screen via the reaction:

$$n + {}^{6}\mathrm{Li} \rightarrow \alpha + {}^{3}\mathrm{H} + 4.78\ \mathrm{MeV} \tag{3.1}$$

The alpha and triton particles produce a scintillation signal in the ZnS(Ag) component of the ^{6}LiF:ZnS(Ag) sheet, before being absorbed by the sheet itself. The scintillation photons from the ZnS(Ag) enter the PVT cube, and are extracted via the optical fibres coupled to the PVT. Thus, the same optical readout system is used for both scintillators forming the cube. At $\mathcal{O}(10)\mu$s, the scintillation time of the ZnS(Ag) is significantly longer than that of EM signals from the PVT at $\mathcal{O}(10)$ns. This provides a means to identify whether a scintillation signal is due to a charged particle interaction or neutron capture in a cube.

For a cube with a single lithium sheet (as is the case with the large scale prototype), the neutron capture efficiency $\epsilon_{Capture}$ can been estimated using the simulation frameworks

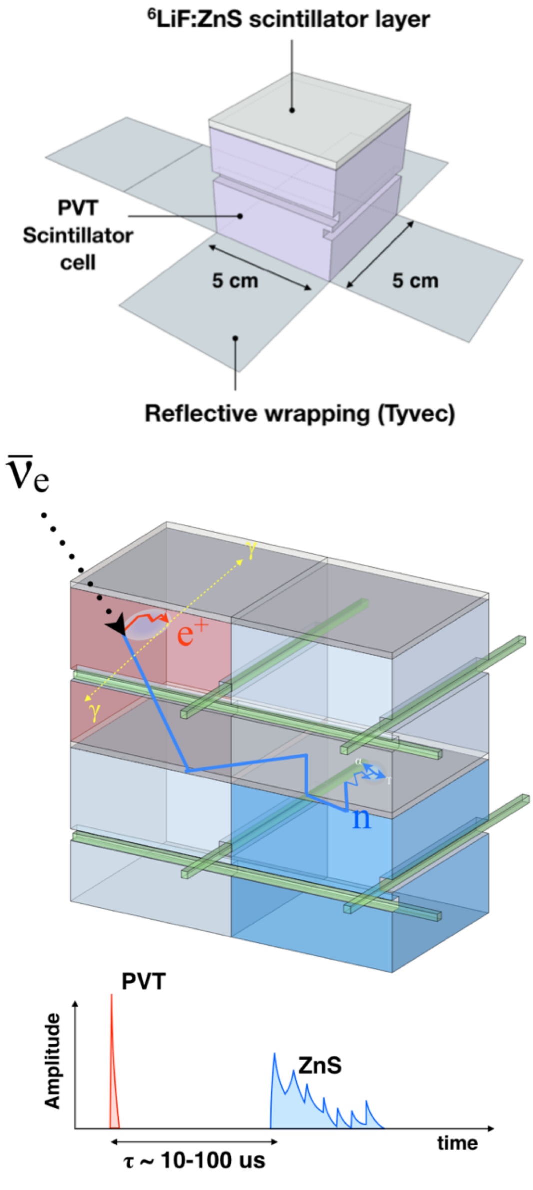

Figure 3.1: Principle of detecting inverse beta decay interactions using SoLid technology. These cubes represent those used for the prototypes - the full detector will have twice the number of optical fibres.

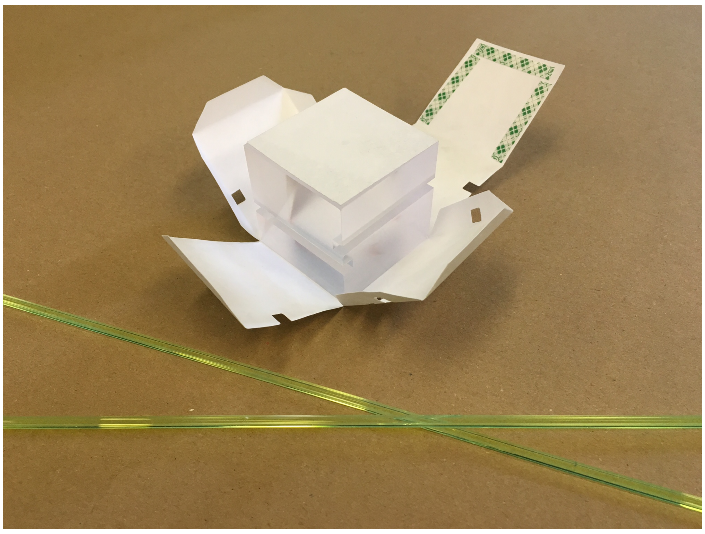

Figure 3.2: Picture of a single cube used for the large scale prototype, with two fibres nearby.

Geant4 and MCNP [49]. The combined result from the frameworks gives:

$$\epsilon_{Capture} = (52 \pm 3)\% \tag{3.2}$$

This is the average of the two frameworks with 1σ errors. For a cube with two sheets of lithium, this increases to around 66%.

3.2.1 Neutrino Signal

As well as providing a means to detect charged particles, the PVT cubes act as a proton rich target for inverse beta decay interactions (IBD). The positron briefly travels through the PVT, giving rise to a sharp EM scintillation signal, before annihilating with an electron from the detector volume, producing two annihilation photons with energy 511 keV. These photons have a mean free path of around 30 cm [4], and therefore mostly deposit energy in cubes away from the positron cubes. This is an advantage of the technology, as the positron energy can be reconstructed independently of the annihilation photons (using only the positron cubes), and thus gamma leakage is not as much of a problem compared to previous technologies [4]. The neutron from the IBD takes on average $\sim 90\,\mu$s to thermalise prior to capture, and typically travels a few cm (< 10 cm [53]) through the PVT that acts as a moderator. This process is sketched in the bottom of figure 3.1. The fact the IBD neutron and positron are highly correlated in both space and time, unlike many large background components in the analysis, gives rise to several new handles to increase the signal to noise ratio of the experiment.

3.2.2 Silicon Photo-multipliers

One end of each optical fibre is connected to a silicon photo-multiplier (SiPM). This type of sensor is especially suitable for single photon counting, whilst at a manageable cost given the highly segmented detector design. The sensors chosen are formed of an array of 60$\times$60 avalanche photo-diodes, each with a quenching resistor, connected in parallel. The active area is 3$\times$3mm^2 (see [40] for the data-sheet). An incident photon may create an electron-hole pair in the photo-diode silicon. Without the application of an external voltage, the pair will likely recombine, giving no detectable signal. Applying a small voltage across the sensor can create a depletion region that is free of charge carriers. Induced electron hole pairs are free to carry charge through this region, although any current induced is too small to be detected. Upon increasing the voltage to above the break-down voltage, the electrons and holes gain enough energy to create a divergent pixel avalanche. The divergence of the avalanche is controlled by the quenching resistor; the voltage across the resistor (due to the current through the pixel) brings the bias voltage across the pixel to below the break down voltage, stopping the avalanche. The output of the SiPM is the total current that flowed through all pixels.

Since the amplitude of the signal depends on the number of pixels experiencing an avalanche, the amplitude of the signals is discretised. The number of pixel avalanches is almost linearly related to the number of incident photons on the sensor, with deviations depending on the voltage setting. As well as controlling whether the sensor is above or below break down voltage, the voltage setting of the sensor has other implications, including:

- **Photon detection efficiency (PDE):** increases with over-voltage. The energy of the electron-hole pairs increases with voltage, which increases the probability of the pairs causing daughter avalanches.
- **Cross talk:** since there are no boundaries in the silicon to separate pixels, a pixel touching a firing pixel may also experience a pixel avalanche, even though no photon was incident on the pixel itself. This will artificially increase the amplitude of the output signal by $\sim 10\%$ on average [63], biasing the energy response and worsening the energy resolution.
- **Quenching:** a photon incident on an already firing pixel will not increase the amplitude

of the signal (as it would if incident on a non-firing pixel). For SoLid, given the expected light yield (see below), this effect is negligible.

- **Dark count rate:** electrons have a significant probability of randomly causing a pixel avalanche due to thermal agitation. This is highly dependant on over-voltage and temperature - at room temperature, the dark count rate is typically in the MHz range, and this approximately halves for each ten degrees of cooling.

The discretised amplitude also depends on the gain (a measure of the number of charge carriers generated by a single photon), which is linearly related to the difference between the set voltage and the break down voltage (i.e the over-voltage). An example waveform from a SiPM is shown in figure 3.3. The distribution of the amplitude of the peaks found in these waveforms is shown in figure 3.4. In these examples, the gain of the sensors was set to $\sim$65 ADC counts per pixel avalanche (i.e. the amplitude of the 1 pixel avalanche signal). The discretised nature of the SiPM signals is observed, and the contributions from the first, second and third pixel avalanche peaks (separated by the gain) are clearly visible.

Gain Equalisation

Since the break down voltage of each SiPM is different [54], online equalisation procedures are required to equalise the amplitude response of all sensors in a detector. This requires finding the break down voltage of each sensor, which can be done by performing a voltage scan. At each voltage, for a given channel, the gain is found from the position of the first and second pixel avalanche contributions in the amplitude spectra of peaks (such as in figure 3.4). The linear relationship between gain and voltage for each channel is then extrapolated to zero, giving the break down voltage of the sensor. This equalisation procedure can be difficult if there are large levels of electronic noise in the system, which smears the position of the first and second pixel avalanche positions in the peak spectra.

3.3 Large Scale Prototype - SM1

Prototype SoLid Module One (SM1) aims to demonstrate the viability of the new technology at the large scale. It was commissioned in early 2015 during a period when the reactor was operational, and just before a one year long shut-down for a reactor refit.

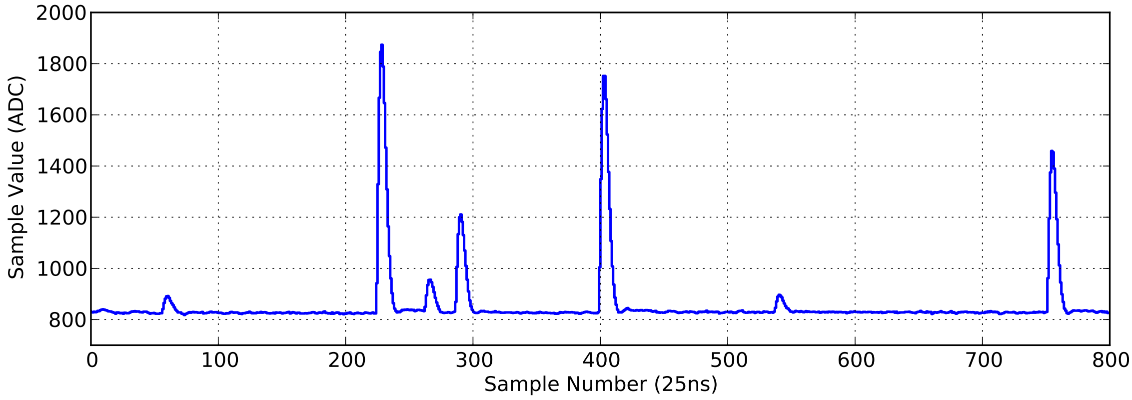

Figure 3.3: Example SiPM waveform from phase 1 test data (randomly triggered). The SiPMs were not attached to any optical fibres, and therefore only dark counts are present.

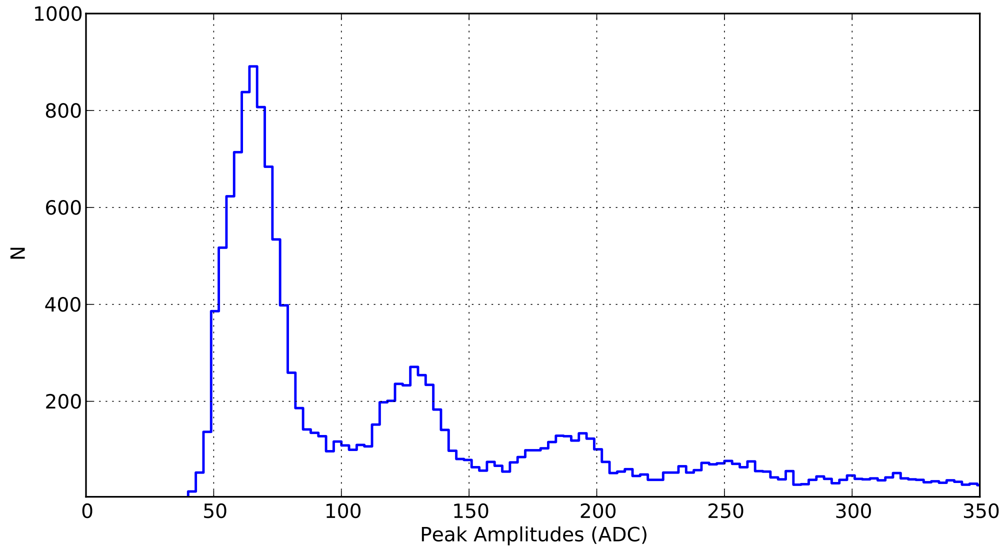

Figure 3.4: Peak amplitude spectrum of phase 1 randomly triggered waveforms. The algorithm used to find peaks in the SiPM waveforms is described in chapter 4.

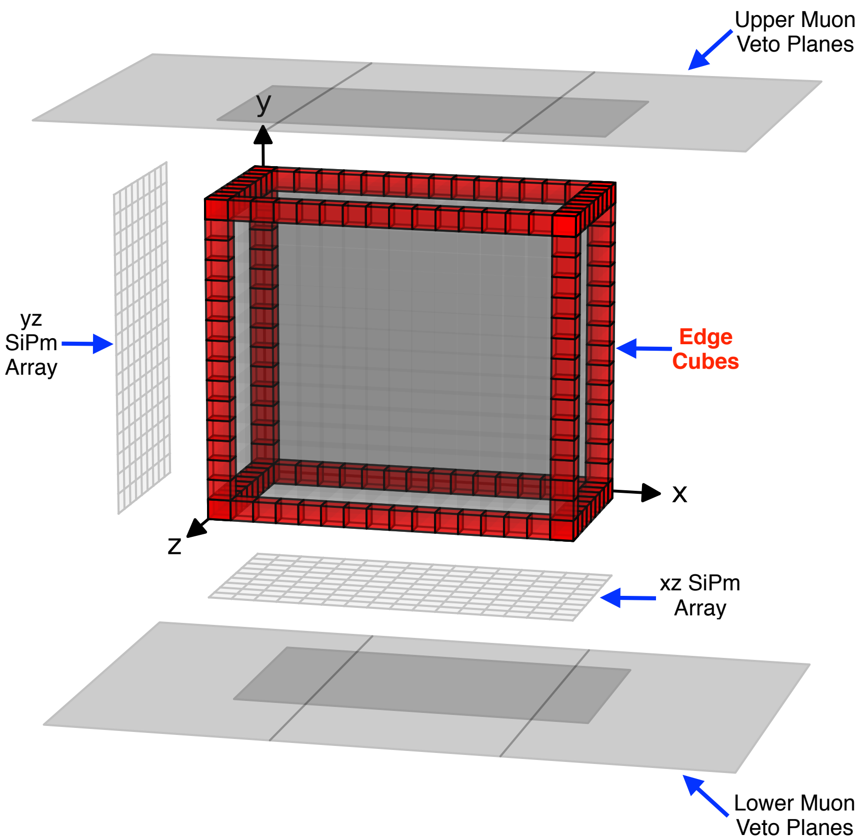

Figure 3.5: The co-ordinate system used for the large scale prototype (SM1). The z axis points away from the reactor. The reactor itself is at negative z. The red cubes show the 'edge' cubes.

3.3.1 Mechanical Configuration

The cubes are stacked in an arrangement of 9 planes of 16×16 cubes. All 2304 are wrapped in 75 g/m^2 thick tyvek, giving a total active mass of 288 kg. Figure 3.5 shows an empty event display illustrating this configuration, and also shows the co-ordinate system used. Each cube is lined with a single sheet of ^{6}LiF:ZnS(Ag), placed on the face of the cube that is perpendicular to the reactor (i.e. in the xy plane), and closest to the reactor. All cubes in a plane are coupled to two fibres (one per xy direction), and each plane has an array of 16 vertical and horizontal fibres to read out the cubes. This is illustrated in figure 3.6, which shows a diagram of a decomposed plane. The picture in figure 3.7 shows a prototype plane being filled with cubes. In each plane, the edges of the cube array are lined with 2 cm high density polyethylene (HDPE) reflector bars, which act to reflect escaping neutrons, and compensate for any mechanical misalignment of the cubes. The array of cubes sit in an aluminium frame, with holes placed such that the fibres sit just outside the frame, to be attached to the SiPMs. The faces of the plane are lined with 1mm sheets of black HDPE for structural integrity. Nine planes are stacked to form the prototype detector, as shown in figure 3.8.

Muon detection panels composed of 2 cm slabs of PVT are also placed above and below prototype detector, to allow for a muon veto system. However, due to the thinness of the planes, and given the large gamma background from the reactor itself, the trigger thresholds are required to be set very high to achieve a manageable data rate. This results in lowering the muon detection efficiency. The operation of these panels has been unsuccessful.

3.3.2 Readout

The cubes are readout using 288 single cladded wavelength shifting fibres. Each fibre is optically coupled to a single silicon photo-multiplier, and a mirror placed at the other end of the fibre to increase the light yield. The sensors are biased with a target overvoltage of 1.5 V, giving a photon detection efficiency $\sim 25\%$, and a cross talk probability of $\sim 10\%$ [63]. Since the prototype had no cooling, the detector operated at room temperature, resulting in a dark count rate of around 0.5 MHz per channel. Due to high levels of electronic noise present, it was difficult to equalise the prototypes sensors online to an accurancy better than $\sim 20\%$ [53], requiring offline calibrations to achieve uniformity in the detector (see chapter 5). Techniques to combat this electronics noise are discussed in chapter 4.

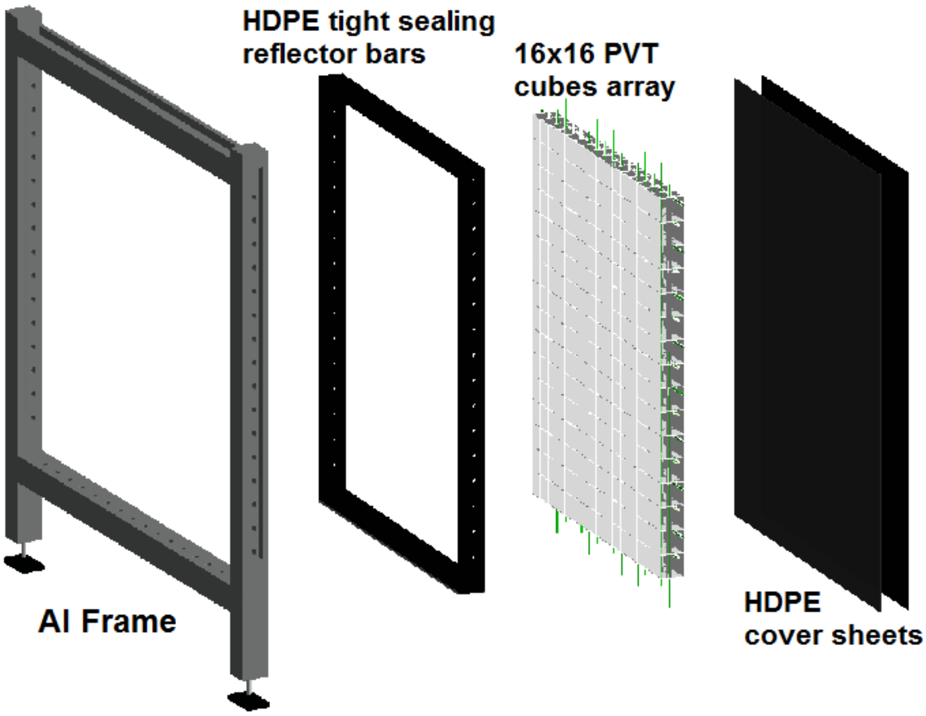

Figure 3.6: A diagram showing a decomposed prototype plane.

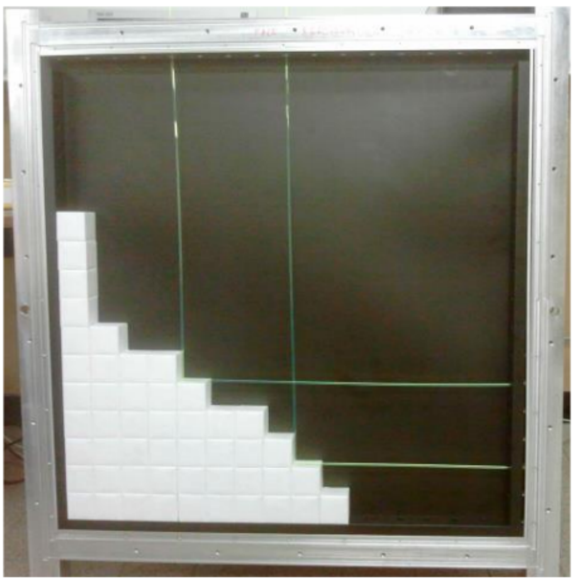

Figure 3.7: A picture of a partially filed plane of cubes, with example optical fibres.

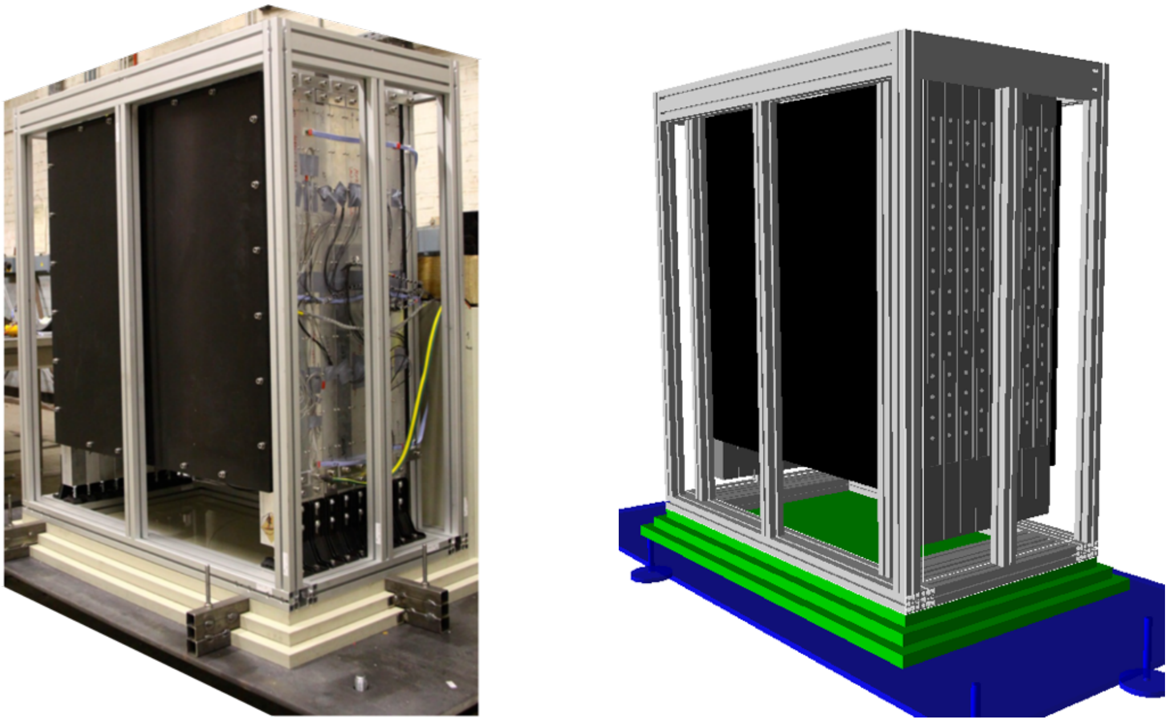

Figure 3.8: Left: Picture of the large scale prototype. Right: Diagram of the prototype from the Geant4 simulation.

Triggers

The prototype detector uses a single trigger configuration to read out all types of particle interactions. Initially, this was a simple threshold-based trigger applied separately to all channels. Upon triggering, the waveform from that channel would be readout. However, it was found during commissioning that due to large levels of electronics noise, and the large reactor gamma background, the trigger thresholds had to be set extremely high to achieve a manageable data rate. In order to readout lower amplitude signals, a co-incidence requirement is also imposed to reduce the data rate. Two channels in the same plane, attached to fibres in orthogonal directions, are required to be above threshold within a time window of 48ns (i.e. 3 waveform samples). Upon triggering, the waveforms (256 samples, 4.096μs) of all channels above threshold are readout. If another trigger arrives during this time, the waveforms readout by the first trigger are truncated. The first 25 samples of the waveform are prior to the trigger time, and the remaining samples are from after the trigger time. Using this trigger, IBD interactions therefore require the neutron and positron to trigger separately.

The use of a co-incidence requirement did achieve a manageable data rate, however the

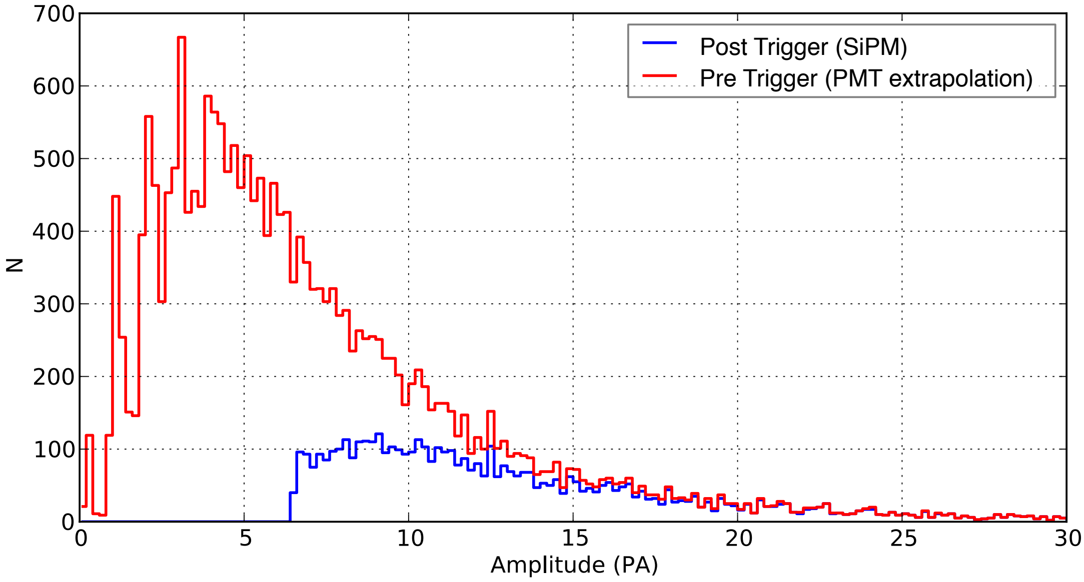

Figure 3.9: Neutron amplitude distribution. The pre-trigger distribution is measured using a PMT and scaled to the light yield of the SiPM. The post-trigger distribution is from a single SiPM, from a lab setup using two SiPMs, which emulates the prototype trigger conditions.

trigger threshold is still set relatively high at ~ 6 SiPM pixel avalanches (PA), and this has had consequences. The neutron amplitude distribution is shown in figure 3.9. It peaks at low amplitudes, and only the tail extends above this trigger amplitude. Further, there is only a weak correlation between the amplitude of ZnS signals from the two readout fibres coupled to a cube. As a result, the trigger efficiency for neutrons is low at around 5% [49]. Similar arguments can be made for other types of interactions, leading to reduced capability to separate signal and background (see chapter 6).

Electronics

Custom analogue boards (one board per plane) have been developed to amplify and digitise the SiPM signals, as well as set the bias voltage for each SiPM channel individually. The analogue signals from these boards are digitised at a rate of 62.5 MS/s (16 ns per waveform sample) with a 14-bit resolution ADC, providing an effect range of 0 to 16 k for waveform sample values. The pedestal of each channel was around 8 k, varying on a per-channel basis. The sampling rate is set fast enough so as to reveal the structure of the ZnS signals. Each pair of analogue boards (i.e two detector planes) was connected to a Gigabit Link Interface Card (GLIB). The trigger logic of the system was implemented on the FPGAs of each GLIB.

Dataset name	Dates	Live time (hrs)
Reactor On	00:00 21 Feb → 08:00 24 Feb	50.9
Reactor Off	08:00 24 Feb → 00:00 12 Mar	577.8
	and 00:00 27 Mar → 12:00 11 Apr	
AmBe calibration	28 Apr	2.4

Table 3.1: Summary of prototype module one data used in this work.

3.3.3 Light Yield

The light yield of the prototype cubes has been measured both in a lab setup and during plane commissioning [63]. A ^{207}Bi source provides electrons with a peak energy of 1 MeV, and can be placed underneath a single cube. It is found that the typical light yield of a prototype cube is 11.2 ± 0.6 (systematic) SiPM pixel avalanches per MeV per fibre (PA/MeV/fibre). The systematic uncertainty has been found by re-assembling the lab setup multiple times, and dominates the statistical uncertainty. The distribution of the number of photons detected from a 1 MeV signal is Poisson distributed, and so the relative energy resolution (taken as the square root of the variance) is $E/\sqrt{N_{photons}}$, or 20% at 1 MeV.

3.3.4 Datasets and Stability

Commissioning finished during a reactor on cycle and 3.5 days prior to a long shut down. During the shutdown period, a long reactor off run was taken for background studies, as well as source calibration runs. The datasets studied in this thesis are summarised in table 3.1. Figure 3.10 shows various trends found in the data, for both the reactor on and off data taking periods. The reconstruction of the various objects is explained in chapter 4. It can be seen that the rate of cosmic muons is stable over this period, suggesting timing of the detector is stable. The same can be said for energy calibrations, which are applied offline using studies of cosmic muons (see chapter 5), and is stable to within 2%. The reactor on-off transition is also clearly visible, and the rate of waveforms read out decreases sharply from around 25 kHz to 15 kHz.

3.4 Phase 1 Detector

3.4.1 Configuration

At the time of writing, Phase 1 of the experiment is being assembled. The initial active mass will be 1.6 tonnes. The detector will be composed of fifty planes, in a 16 × 16 configuration

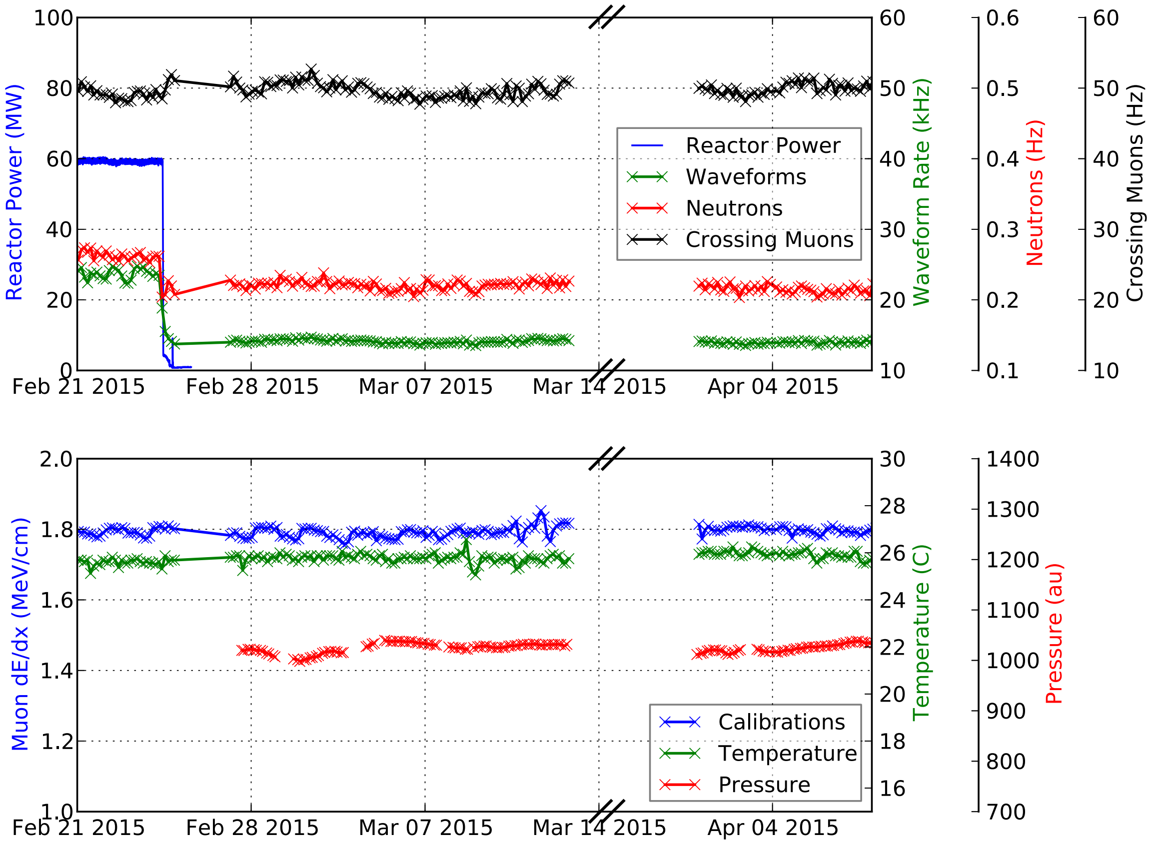

Figure 3.10: Data trends for prototype module one data taking. The second half of March has been removed due to various environmental tests conducted in the BR2 building itself.

like the prototype. The shape (i.e. ratio of width/height to depth) has been optimised taking into account the sensitivity of performing the oscillation search, as well as mechanical and electronics constraints. The components of the design have been optimised based on experience from the large scale prototype, including:

- **Light yield**: the number of optical fibres has been doubled - each cube is read out using 4 optical fibres (two per direction). Combined with thicker tyvek and double cladded fibres, the light yield is expected to increase to 14.0 ± 0.5 PA/MeV/fibre, improving the energy resolution from 20% to 16% at 1 MeV.

- **Neutron trigger**: a dedicated neutron trigger will be used to identify ZnS signals in the electronics trigger. The purity of this trigger will be significantly higher than that of a threshold-based trigger, allowing for more manageable data rates, and increasing the overall neutron detection efficiency.

- **Cooling**: the entire detector will be placed in a temperature controlled environment inside a shipping container. The target temperature is 5 °C, significantly lowering the dark count rate of the SiPMs, which should increase the trigger purity further. The cooled environment will also stabilise the gain calibration of the SiPMs.

- **Two sheets ^{6}LiF:ZnS(Ag)**: the number of lithium sheets will be doubled compared to the prototype, increasing the neutron capture efficiency from 54% to 66%.

- **Passive shielding**: four walls of 50cm water bricks will be built around the detector, as well as ∼50cm of HPDE placed on the roof of the container, for increased shielding from fast neutrons.

Many of these improvements, such as the choice of neutron trigger algorithm and shielding effects, are discussed in more detail in chapter 7. A drawing of the upgraded detector is shown in figure 3.11. The detector is composed of 5 modules, each composed of ten planes. Each plane has a dedicated electronics ‘cassette’, to house various electronics components. Each module is powered separately, and has a dedicated cooling system composed of a chiller fan (pushing air downward), and heat exchanger.

3.4.2 Calibration

Figure 3.11 also shows the calibration robot CROSS, which is placed above the cubes. Periodically, source calibration runs will be performed. The modules will be mechanically separated, and a robot is used to place a source at any position in the gap between the two modules. It is foreseen that 9 positions will be scanned. This should provide an absolute neutron efficiency calibration accurate to $\sim 3\%$, and absolute energy scale calibration accurate to $\sim 1\%$ [55].

3.5 Belgian Research Reactor 2 - BR2

Phase 1 of SoLid, like the large scale prototype, will study the Belgian Research Reactor 2 (BR2) at the SCK•CEK research centre, Mol. BR2 operates at powers ranging from 50 MW to 80 MW [67], with the reactor typically producing power for half a year in periods of around a month. The rest of the time, the reactor is inactive. This reactor is highly suitable for SoLid, since:

- Being a research reactor, the core is compact with a diameter $\sim$ 0.5m. This is small relative to the baseline being studied, reducing the uncertainty in measuring the baseline. This is an advantage compared to power reactors, which can have cores 10$\times$ wider.
- The fuel is $\sim$ 95% highly enriched ^{235}U [67]; a measurement of the neutrino energy spectrum can help resolve the source of the 5 MeV distortion.
- The layout of the reactor hall allows for detectors to be placed very close to the reactor core (as close as 5 m standoff).

There are many possible production decays resulting in neutrinos inside nuclear reactors. For ^{235}U, the most likely decay chain is:

$$\mathrm{n} + {}^{235}_{92}\mathrm{U} \rightarrow {}^{141}_{56}\mathrm{Ba} + {}^{92}_{36}\mathrm{Kr} + 3\mathrm{n} \tag{3.3}$$

The ${}^{141}_{56}$Ba (Barium) and ${}^{92}_{36}$Kr (Krypton) atoms are unstable, and are most likely to decay to ${}^{94}_{40}$Zr (Zirconium) and ${}^{140}_{58}$Ce (Cerium) respectively via beta decay, producing six neutrinos.

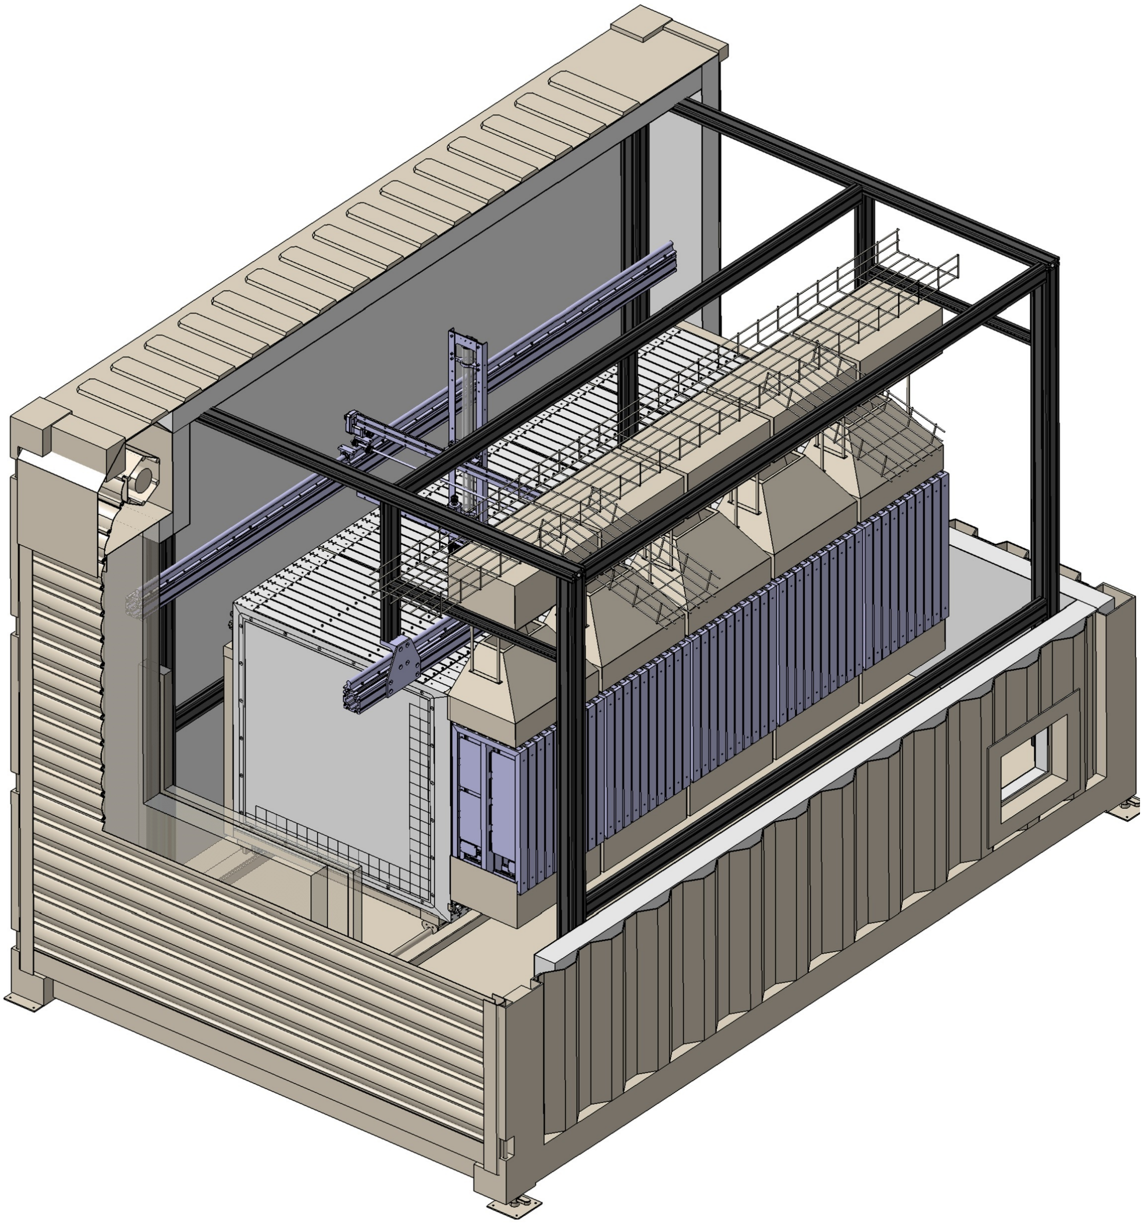

Figure 3.11: Drawing of Phase 1 of the SoLid experiment. The CROSS calibration robot hovers above the modules. The entire detector is placed inside a temperature stable shipping container, and will be cooled to around 5°C.

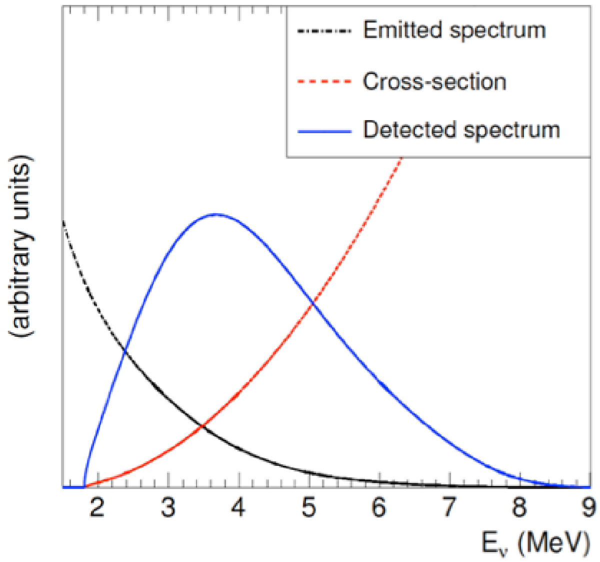

Figure 3.12: Example energy spectrum of emitted reactor neutrinos, and the detected neutrino energy spectrum. The cross section increases with neutrino energy, and so the detectable neutrino energy is peaked.

An example energy distribution of reactor neutrinos (similar to that expected at SoLid) is shown in figure 3.12, as well as the cross-section of IBD interactions. When combined, it can be seen that the energy distribution the neutrinos that interact via IBD is peaked around 3.5 MeV.

A schematic of SoLid and the BR2 reactor is shown in figure 3.13.

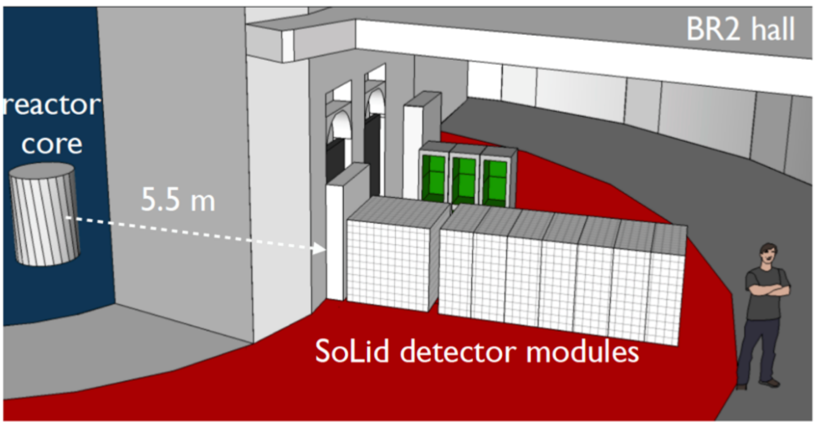

Figure 3.13: Schematic of SoLid at the BR2 reactor.

3.6 Conclusions

The SoLid collaboration have developed a new technology to tackle many of the challenges associated with performing very short baseline reactor neutrino experiments. This features a highly segmented design, using 5cm^3 cubes of composite scintillator, PVT, to detect EM signals (i.e. charged particles and photons); and ^{6}LiF:ZnS(Ag), to detect neutrons. Light is emitted from these two scintillators at very different time-scales, allowing the type of signal (EM or ZnS) to be identified. A two dimensional array of optical fibres is used to read out each cube, providing excellent position resolution in 3D. The segmentation allows the positron energy to be reconstructed independent of the annihilation gammas. Silicon photo-multipliers are attached to the ends of the fibres, and allow single photon counting. A large scale 288kg prototype using the technology, with around 20% the mass of that planned for phase 1 of the experiment, was deployed in early 2015. The analysis of the data taken with this prototype is presented in the following chapters. Phase 1 of the experiment is currently under construction. Its design is the result of an optimisation of the prototypes: the energy resolution has been significantly improved, due to material choices and doubling the number of readout channels; a dedicated trigger configuration is used to perform neutron identification at the electronics level; and a 50 cm water wall will be built around the detector, with a 50 cm HDPE plastic roof, to provide passive shielding from cosmic radiation.

The prototype (like SoLid phase 1) uses the the Belgian Research Reactor 2 (BR2) in Belgium at SCK•CEN as a neutrino source. This reactor is especially suitable, given its compact reactor core, and suitable environmental conditions for performing the experiment. The reactor itself uses highly enriched Uranium-235 as fuel, and measurements of the core's neutrino energy spectrum can help resolve the source of the 5 MeV distortion reported by previous reactor experiments.

Chapter 4

Event Reconstruction

4.1 Introduction

SoLid detectors output data in the form of waveforms registered by silicon photo-multipliers. Conversely, the input data of an oscillation search for sterile neutrinos are the location and energy of neutrino IBD interactions. The following three chapters will bridge these two stages of the experiment, by performing reconstruction of neutrino candidates from raw detector data. Since prototype SoLid module one was the first large scale SoLid detector, many of the following reconstruction algorithms have been developed for the first time, and are used by the analyses presented in later chapters. As such, preference is often given to simpler algorithms, whose features can be understood more easily. Cases of ongoing work to improve them for the next phase of experiment are briefly described.

The chapter is structured as follows: the remainder of the introduction provides some definitions and time-scales referenced in this chapter, and briefly outlines the event reconstruction chain. The rest of the sections are dedicated to the steps in this chain, including event finding, the application of calibrations, and performing particle identification to categorise events. The implementation of the reconstruction in software is also briefly described.

I am grateful to Leonidas Kalousis who implemented the FFT correction inside the SoLid analysis framework (Saffron1).

4.1.1 Event Definition and Time-scales

A SoLid event is defined as a set of co-incident signals, from multiple detector channels, at the scale of the timing resolution of the detector - i.e. $\mathcal{O}(10)$ns. This is often referred to as the fast time scale of SoLid, and therefore includes any interactions in the detector that give rise to simultaneous signals between channels, such as positrons, neutrons, and muons. This chapter focuses on reconstruction of these events at this fast time scale; chapter 6 deals with reconstruction at the much longer neutron thermalisation time scale $\mathcal{O}(100)$ μs, such as forming inverse beta decay candidates from neutron and positron events.

4.1.2 Event Categories

This chapter will outline the reconstruction of the following types of event:

- **Neutron candidates**: scintillation photons from the ZnS. These signals have relatively low amplitude, and have a characteristically different waveform shape compared to EM signals. Effective (i.e. efficient and pure) separation from the large number of EM signals is required to reduce backgrounds in later analyses.

- **Muon candidates**: very large energy depositions across many channels. Since muon spallation can lead to neutrino backgrounds, these must be efficiently identified. Additionally, muons with a long path-length through the detector can be tracked, and can be useful for calibrations.

- **IBD Prompt candidates**: a relatively large EM energy deposition in a small cluster of touching cubes, in fast co-incidence with two annihilation photons. Accurate energy reconstruction of the positron is critical for the sterile neutrino search. Given the detector is not optimised to separate positrons from other sources of EM signals, this category of event is highly contaminated by indistinguishable backgrounds, such as reactor photons and proton recoils.

To separate events into these categories, various features are found for each event, such as topological shape and energy deposited. Selection criteria based on these features are applied later to categorise events.

4.1.3 Different Detector Configurations

Unless otherwise stated, the following results use data or simulations of the most recent large scale SoLid detector prototype: SoLid Module One (SM1). While methods presented are optimised for this large scale prototype, they largely also apply to the previous prototypes and the future full-scale SoLid detector.

4.1.4 Brief Overview of Reconstruction Chain

The steps in the event reconstruction chain can be summarised as:

- Waveform cleaning: fluctuations in the waveforms, such as discontinuities between sequential samples and periodic additional frequencies due to electronic pick-up, are removed to increase uniformity of SiPM response across channels.
- Event finding:
  - Peak finding: extracting signal peaks from waveforms.
  - Peak time clustering: grouping peaks from multiple SiPM channels using timing information to form events.
- Calibration application: channel and attenuation corrections applied.
- Particle identification: selections applied based on event features to separate into categories.

Each step is now described in more detail.

4.2 Event Finding

4.2.1 Waveform Cleaning

The additional fluctuations in the waveforms are non-uniform across channels. Such non-uniformities can lead to many difficulties at later stages in analyses, such as varying efficiency and purity across the detector channels. To reduce this effect, several waveform cleaning techniques are applied. These fluctuations can be split into two categories: glitch fixing and periodic pick-up.

Glitch Fixing

Discontinuities or 'glitches' are present in approximately 10% of the digitised waveforms - see figure 4.1 for an example. The cause of these glitches is currently unknown, but thought to occur during the digitisation process. This leads to 1-3 waveform samples to be erroneously higher or lower than the surrounding waveform samples. The values of these glitch samples can be both positive and negative, depending on the value of the sample prior to the digitisation.

These glitches can have several consequences. They often lead to a large number of triggers that are otherwise not present, which saturate the DAQ data rate. Further, methods for energy reconstruction and neutron identification both rely on the integral of these waveforms, which are often erroneous as a result of these glitches.

Glitches can be found by considering the derivative between sequential waveform samples. Since the glitches are found at the extremes of the ADC range, this gives a high derivative at the rising and falling edge of the glitch. As glitches can be both positive and negative, the absolute value of the derivative is used in practice. The distribution of these derivatives is shown in the top of figure 4.2, where the peak near 8000 ADC/sample corresponds to glitches. The peak sits on a background that is over two orders of magnitude smaller, and these entries correspond to large energy peaks, most likely muons. Since this background is also small compared to the whole population (less than one waveform sample per million would be wrongly corrected), the absolute derivative is found to be a sufficient glitch tagger. Waveform samples are corrected if the absolute value of the derivative exceeds 7500 ADC/sample.

Glitches are corrected by linearly interpolating the values of the waveform samples either side of the glitch. An example of the resulting waveform is also shown in figure 4.1.

Improvements

This method can be extended to improve glitch separation by also including the falling edge of the glitch, in addition (i.e. summed) to the rising edge, as also shown in figure 4.2. Note that this currently presumes glitches are a single sample wide - the peak still present at approximately 8000 ADC/sample corresponds to glitches width > 1. This technique is being considered for future productions.

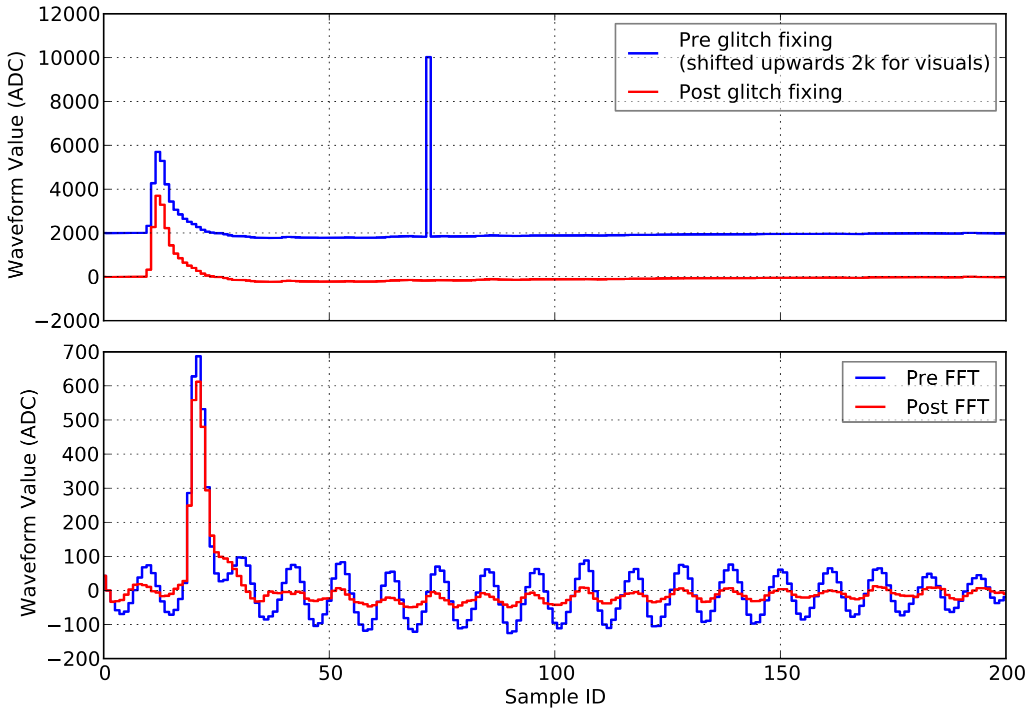

Figure 4.1: Examples of the two types of offline correction applied to waveforms. Top: glitches that are discontinuities in the waveforms. Bottom: periodic pickup, corrected by the use of a FFT..

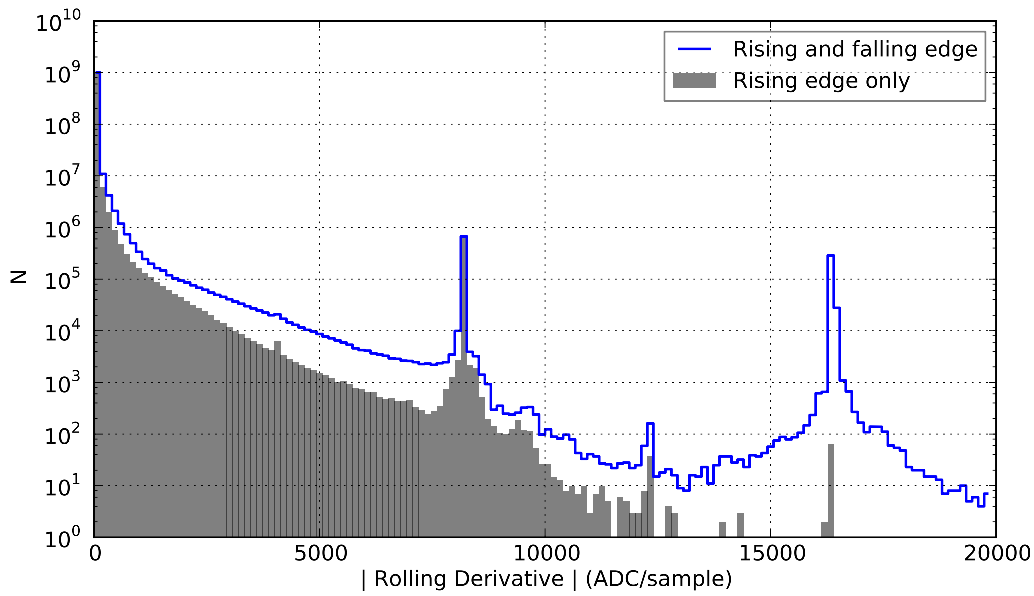

Figure 4.2: Comparison of different algorithms used to identify glitches in waveforms. These algorithms use the absolute derivative between sequential waveform samples (one entry per sample).

Periodic Pick-up

Oscillatory signals in the waveforms are caused by the amplifier circuit and switching noise from the high voltage converters. This source of effect is especially non-uniform across channels, and can also have several consequences. During data taking, signals on the SiPMs that are in co-incidence with a peak of the pick-up are more likely to trigger. Conversely, signals in co-incident with a trough of the pick-up have reduced likelihood of causing a trigger, thus resulting in a time dependent trigger efficiency. Further, the integral of the waveform is dependent on the phase of the pick-up relative to the peak position in the waveforms.

Since this pick-up is periodic in nature, a Fast Fourier Transform can be used to identify it's frequencies. The identified frequencies are used to construct a template waveform, which is subtracted from the original waveform (with the phase extracted using data). The RMS of waveform samples from random triggers can be used as a measure of the fluctuations of a particular channel, and the distribution of these RMS values before and after the application of the FFT are shown in figure 4.3. For reference, the gain of all channels was set such that a single pixel avalanche peak has an amplitude of 25 ADC counts. It can be seen that there is significant reduction in the overall level of fluctuations, as well as increased uniformity. An example of the correction is shown in the bottom of figure 4.1.

The next step in event reconstruction involves finding signal peaks in these cleaned waveforms. The effects of cleaning the waveforms is further demonstrated below.

4.2.2 Peak Finding

Waveforms read-out by the threshold trigger should be one of two shapes (post cleaning): EM, and ZnS. EM waveforms are formed of one SiPM pulse, that is consistently in the same position relative to the start of the waveform - see figure 4.4 for an example. For neutron events, the waveform is formed of multiple SiPM pulses, as seen in figure 4.4. The secondary pulse positions vary considerably. A peak finding stage is introduced to find the position of the SiPM pulses, and extract their features.

Peak Finding Algorithm

A simple peak finding algorithm is used to locate local maxima in the waveforms. The algorithm loops over all samples in the waveform, and a peak is found if the following two

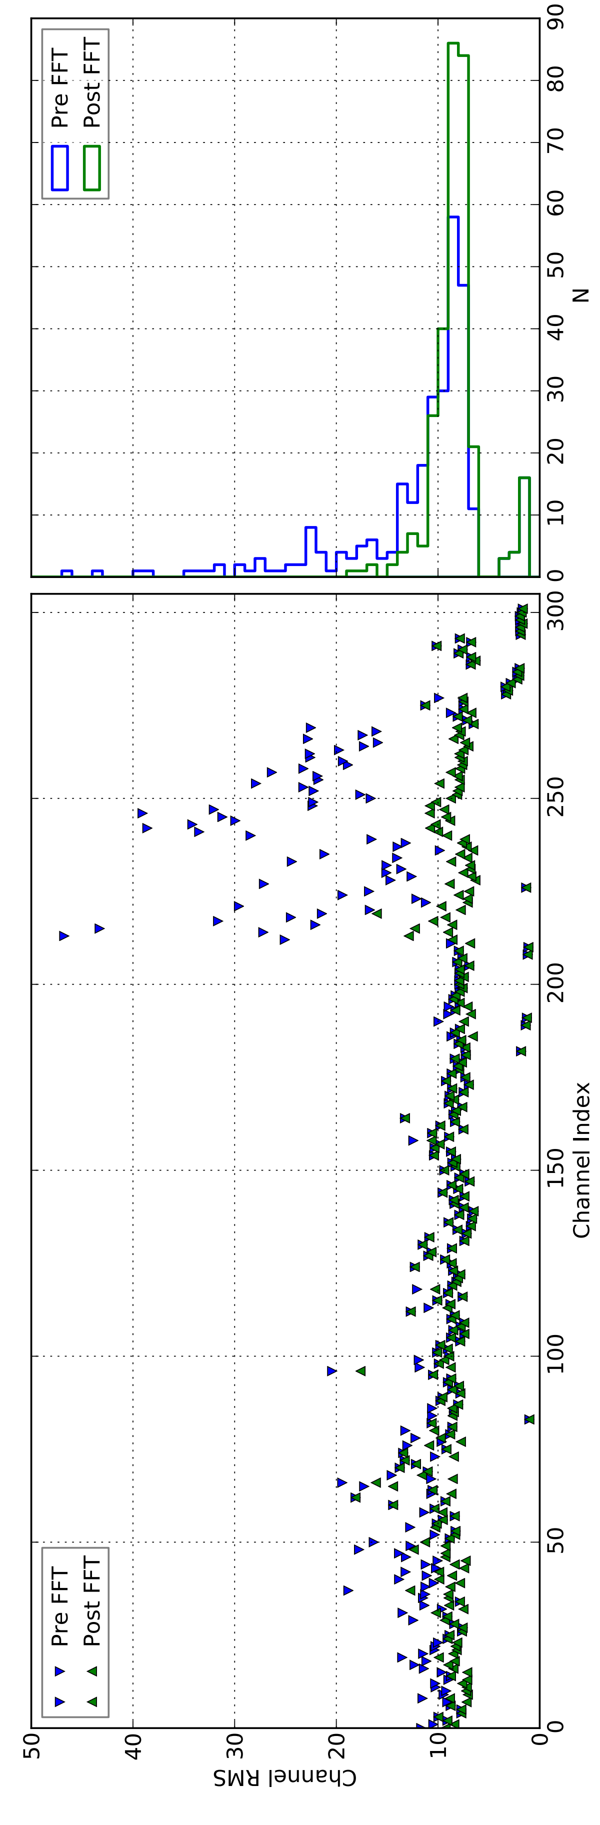

Figure 4.3: Comparison of channel waveform RMS (used as a measure of wave quality), measured using random triggers, before and after the application of a Fast Fourier Transform.

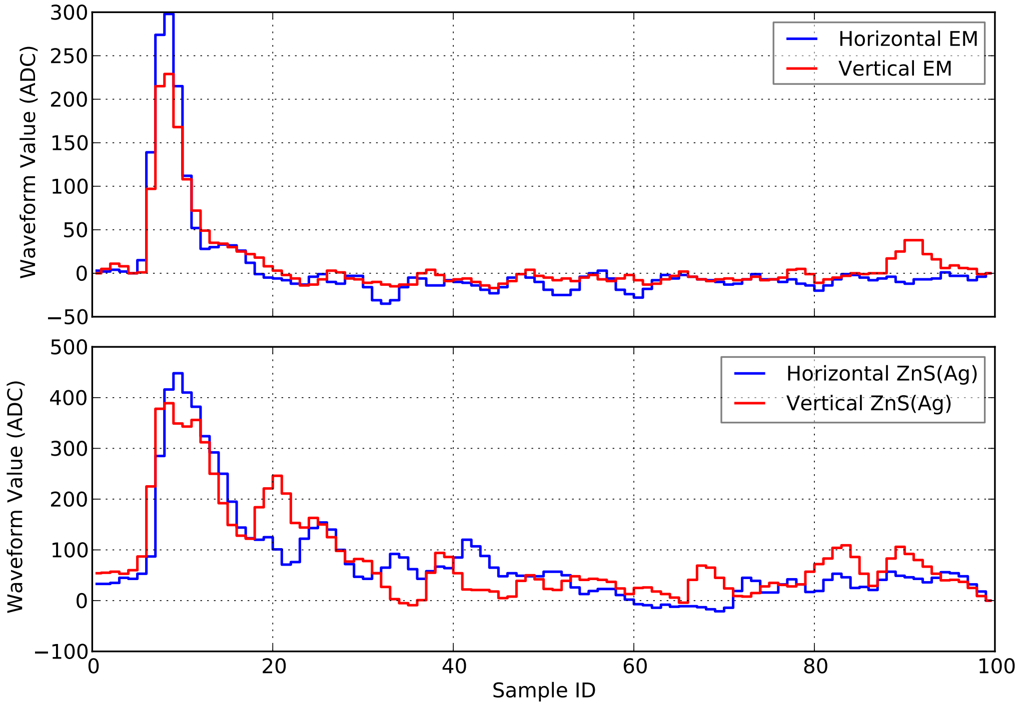

Figure 4.4: Examples of EM candidate (top) and ZnS(Ag) candidate waveforms (bottom). These particular events are for an IBD candidate from the reactor on dataset.

conditions are fulfilled:

- The value of the sample is greater than both its neighbours.
- The value of the sample is above a pre-set threshold. This is used to reduce the large number of background peaks due to electronics fluctuations in the waveforms. The current default choice of threshold is that used by the trigger.

Peak Feature Extraction

The following features are extracted for each peak:

- Time: taken as the time-stamp of the local maximum sample.
- Amplitude: taken as the value of the local maximum sample.
- Integral: the sum of sample values for a fixed range over the peak (including the rising edge of the peak). The range is optimised below for ZnS-EM separation, set at 40 samples.

In addition, the following features are also calculated for each peak from these extracted features:

- Energy: assigned by linearly scaling the integral of the waveform. This is the topic of chapter 5, which details how this scale is determined using cosmic muons for each cube individually.
- I/A: the ratio of integral to amplitude. As will be seen, this is an effective neutron identification parameter.

The 2D distribution of integral vs. amplitude is shown in figure 4.5 for part of the AmBe dataset. This dataset is shown to enhance the neutron contribution for demonstration. Many effects can be observed from these plots. The relationship between amplitude and integral is highly linear. From the left plot, there is a separation of two bands at high amplitudes. This correlates with type of electronics amplifier - two types of channel amplifier were used during the construction of SM1, resulting in two distinct waveform shapes. The effect of these differences will be removed during calibrations outlined below. At high amplitudes, these bands are mostly populated with muon events. At lower amplitudes, an enormous population is observed, corresponding to the low energy EM radiation from the source and the environment. A population at larger values of I/A can also be observed, corresponding to neutron events from the source. The amplitudes of all peaks terminate above 8k (pre-calibration), and this is due to ADC saturation. Outlier populations are still present post waveform corrections, although they are significantly reduced and are away from the neutron population.

4.2.3 Peak Time Clustering

Particle interactions in the detector gives rise to SiPM pulses on multiple channels simultaneously. Reconstructed peaks are clustered in time at the fast time scale to form events, which represent the particle interactions. This gives access to new kinds of variables, including topological. To find correlations between channels, consider the time separation between peaks found on all SiPM channels. This distribution is shown in figure 4.6, and shows the time difference between sequential peaks found from all channels (ordered in time). The fast time-scale of SoLid is clearly visible as an exponential component between 0 and 50 ns. It can be clearly seen that clusters of peaks are efficiently contained within <50 ns, with contaminants from the slower time-scale background at the level of around one per thousand.

A standard one-dimensional seeded clustering algorithm is used to group these peaks into events. The algorithm performs as follows:

- Loop over all peaks (pre-ordered in time):
  - Use the first peak not tagged as belonging to a cluster as a seed.
  - Loop over the remaining peaks.
    - If the time difference between a peak and the seed is less than a pre-specified time window, $\Delta t_{Clustering}$, add to the cluster.
  - Once all peaks have been added, extract the event features.

The value of $\Delta t_{Clustering}$ is determined by the time resolution of the detector, since peaks forming events are simultaneous. The choice of this window should not be made too small as this could lead to partially reconstructed events (AKA 'event clipping'), but also made not

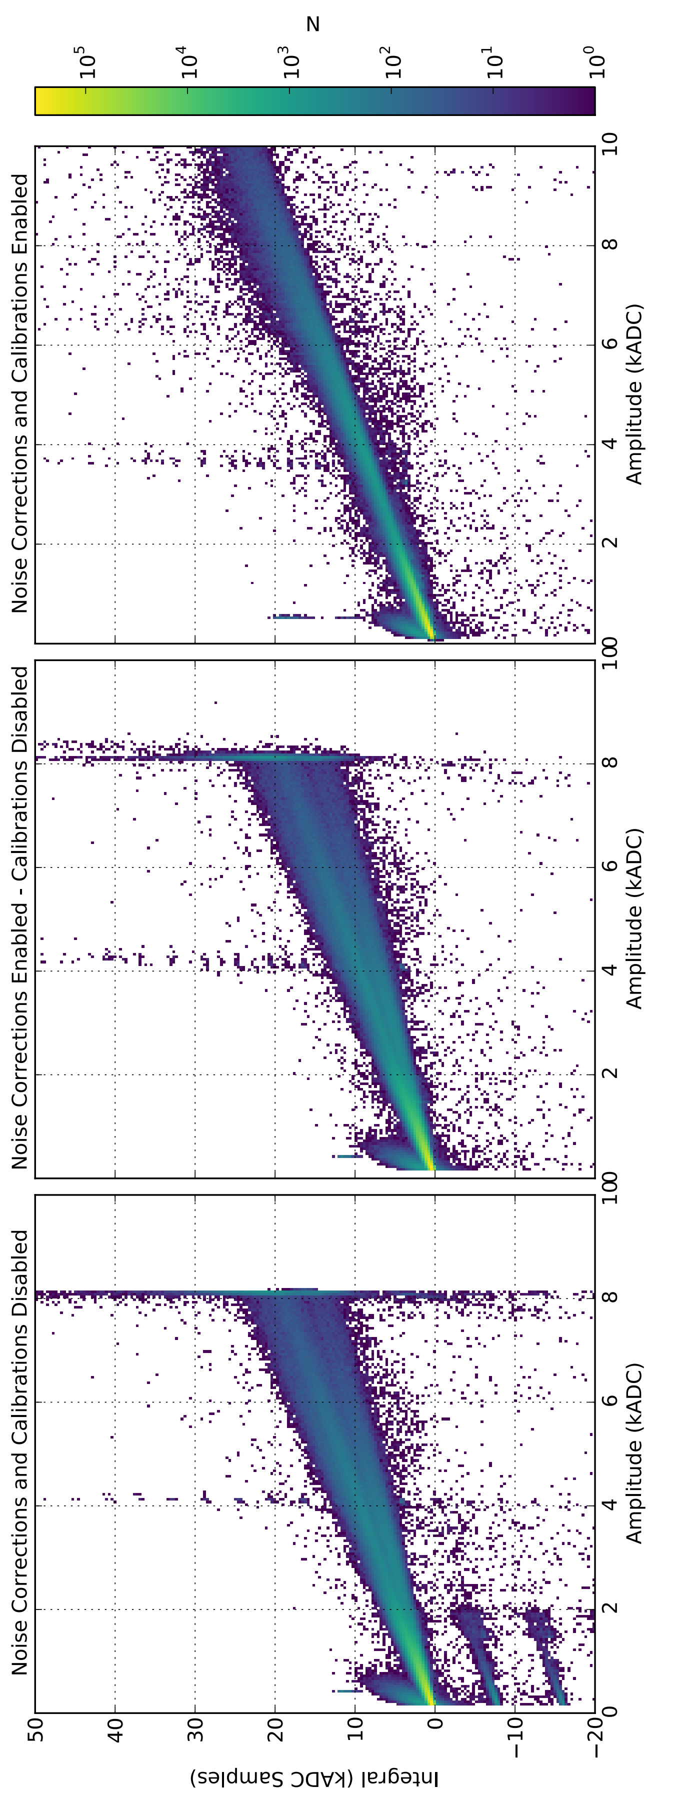

Figure 4.5: Distribution of integral vs. amplitude for peaks identified in waveforms. Left: without waveform corrections and muon calibrations. Middle: with waveform corrections, but without muon calibrations. Right: with both waveform corrections and muon calibrations.

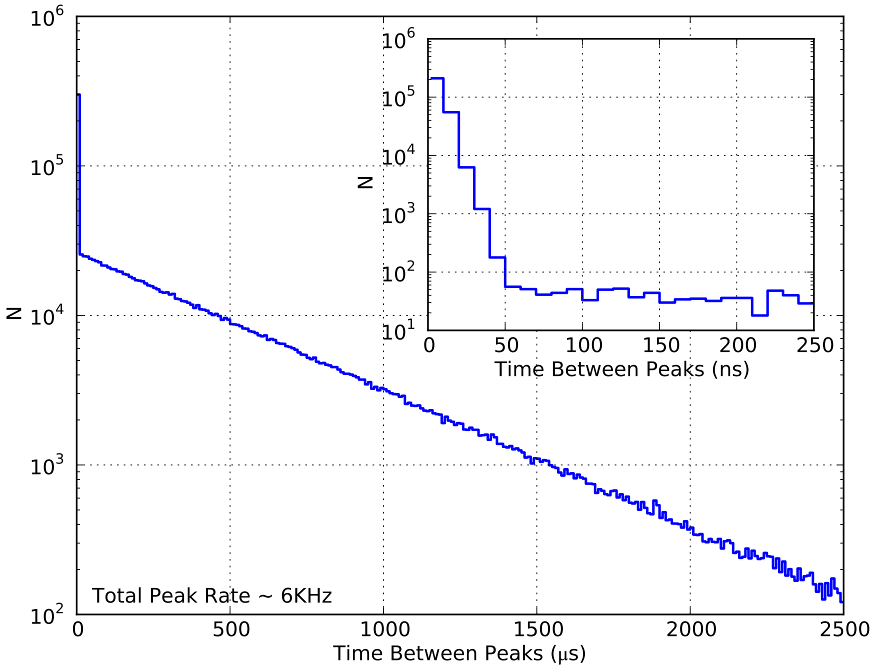

Figure 4.6: Time between reconstructed waveform peaks, found on all channels, and ordered in time. This is for reactor off data.

too large to avoid a high rate of events merging. As observed in figure 4.6, a choice of 50 ns provides efficient clustering with contaminants at a negligible rate.

It is important to ensure that this choice is also appropriate for the two different shapes of waveform signal: EM and ZnS. To demonstrate this, figure 4.7 shows the cumulative distribution of the time difference between horizontal and vertical peaks for both EM and ZnS cube events (i.e. events where only one vertical fibre is found in co-incidence with one horizontal fibre). As can be seen, the ZnS signal is slower, but still well contained to above $> 99\%$ for a time window of 50 ns for both types of signal. It should be noted that this is post trigger selection, and the co-incidence trigger uses a time window of 48 ns; this high relative efficiency is not surprising.

Event Feature Extraction

Events provide a rich set of features (i.e. variable) to separate them into different categories. The features extracted for each event include:

- Time stamp: taken as the average of the forming peaks.

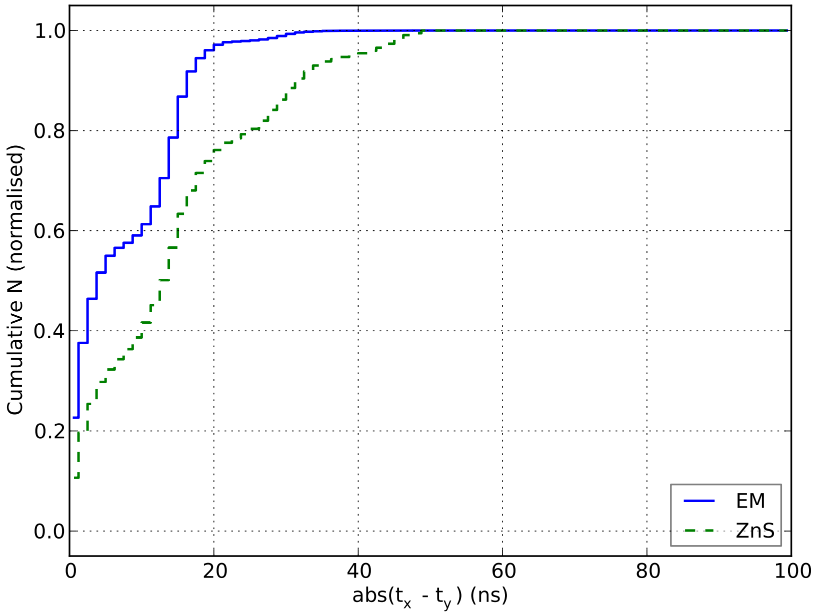

Figure 4.7: Time between vertical and horizontal peaks for EM and neutron events that are localised to a single cube.

- Position (x, y, z): taken as the average the forming peaks.
- Energy: the average of the energy deposits on each SiPM array (i.e. vertical SiPMs and horizontal SiPMs). The energy deposit on a single array equals the sum of all peaks forming the event belonging to that array.
- Topological:
  - Size xy: number of forming peaks found on channels attached to horizontal fibres.
  - Size xz: number of forming peaks found on channels attached to vertical fibres.
  - Width (x, y, z): taken as the maximum width of the event projections in each direction.
  - Volume: the product of the widths (see figure 4.9).

For an example, many of these features have been marked on the event display in figure 4.8, which shows a muon candidate from reactor off data. These variables are now used to classify different kinds of events, and their distributions are shown throughout this and the next chapter.

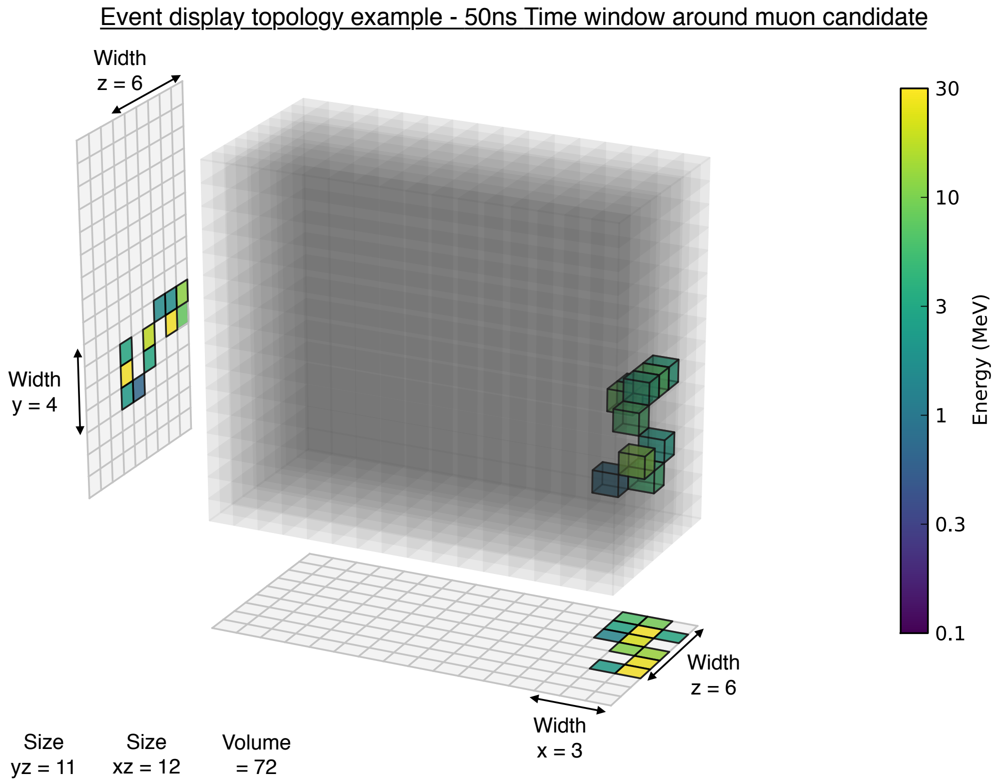

Figure 4.8: Event display of a muon candidate from data, with extracted features annotated. For explanations about reading event displays, see section 4.5.4. The cube finding algorithm used to highlight cubes in the event display is known not to be highly efficient for muon events, and are just included for aesthetic example.

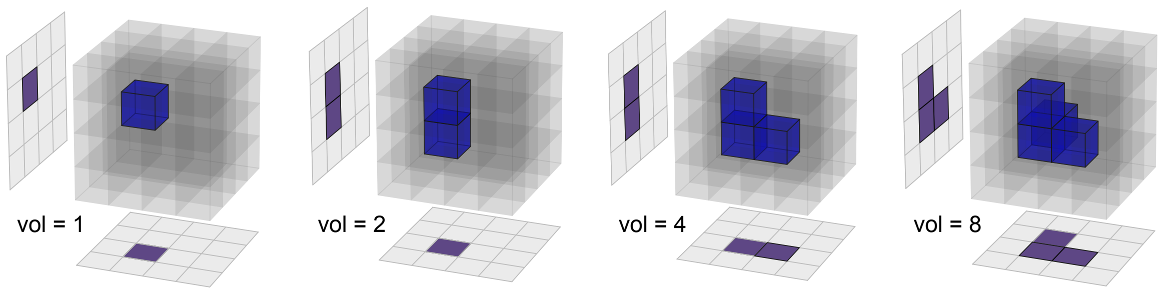

Figure 4.9: Example toy events demonstrating different values of the volume variable, which is a measure of localisation. A smaller detector configuration is shown for visual clarity.

4.3 Particle Identification

Events are classified according to their particle species. For later analyses, it is crucial to have effective separation for neutrons, IBD prompts and muons. These are now described in turn.

4.3.1 Neutrons

As demonstrated by the examples above in figure 4.4, neutrons have a characteristically different waveform compared to electromagnetic signals, where light is sporadically emitted over an extended period compared to EM signals. This gives rise to a larger integral of the waveform relative to its initial amplitude, as also demonstrated by the population at increased I/A shown in figure 4.5. This variable is investigated as a discriminating parameter. Since neutron events give rise to two waveforms from a single cube (one per fibre), there are two values of I/A for each neutron event (selecting single cube events only). The correlation between them is shown in figure 4.10 for the AmBe dataset. Both the EM and ZnS populations can be observed. In practice, classification criteria uses the sum of the two values: $n_{PID} = I_x/A_x + I_y/A_y$. This is equivalent to placing a diagonal cut, as also marked on figure 4.10.

The projection along this diagonal is shown in figure 4.11 for multiple datasets. The two populations are well separated, with the EM population around 5 orders of magnitude more populous. The two neutron source runs using AmBe clearly have larger populations at higher n_{PID}, which is not observed for the EM source run using ^{60}Co, demonstrating that this population is indeed dominated by neutrons. The reactor on run is consistent with the reactor off run for the neutron region, with the larger EM population due to the gamma radiation from the reactor.

The default selection for an event to be categorised as a neutron is: $n_{PID} > 10$. The efficiency of this selection is evaluated by modelling the neutron population as a Gaussian distribution - see figure 4.12. It is found that this default selection, for identification, has an efficiency of $98.3 \pm 0.3\%$. The errors are taken from those of the fit, and these errors dominate the systematic changes observed when varying the fit range. This fact, combined with a reduced χ^2_{dof} of 50.08/47, suggest that a Gaussian approximation is appropriate for the neutron population.

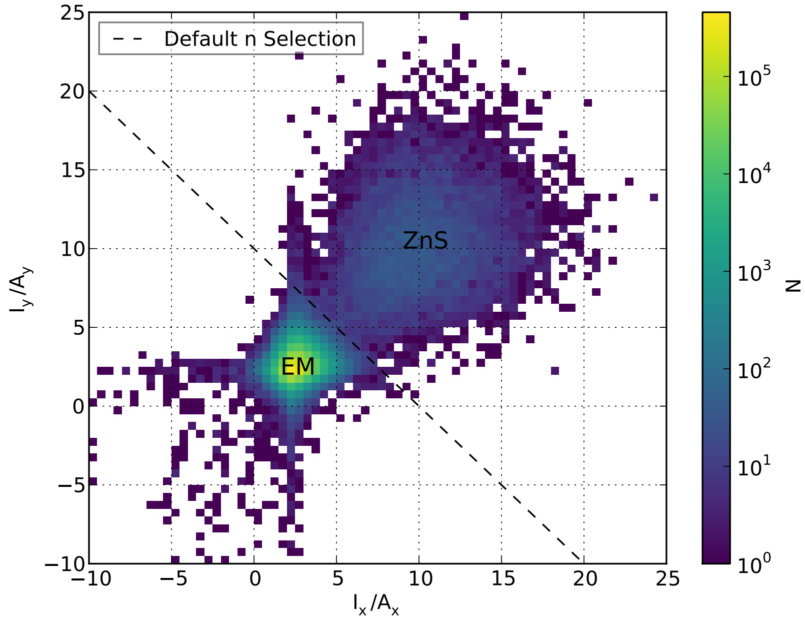

Figure 4.10: Correlation between the two I/A neutron discrimination parameters for single cube events (one value of I/A per fibre direction). This uses the AmBe dataset. The default neutron selection is above the marked diagonal.

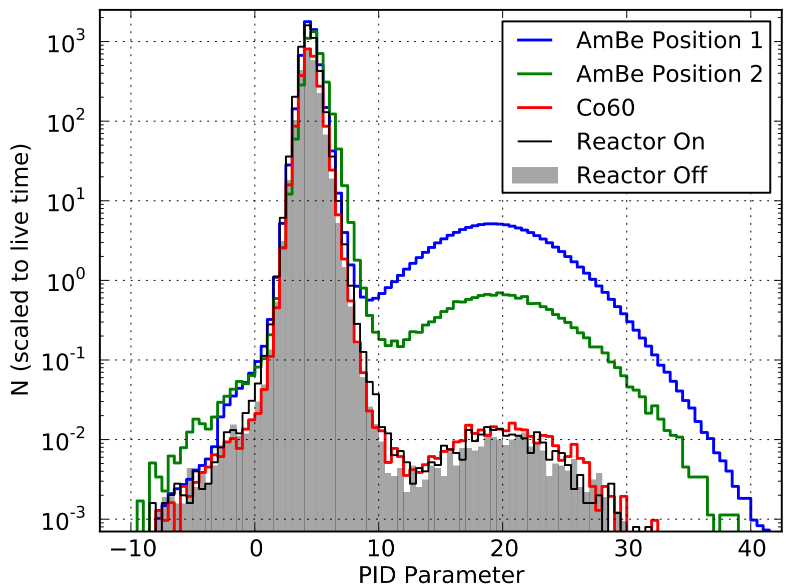

Figure 4.11: Neutron PID parameter for multiple datasets, using the sum of the two ratios of I/A for single cube events.

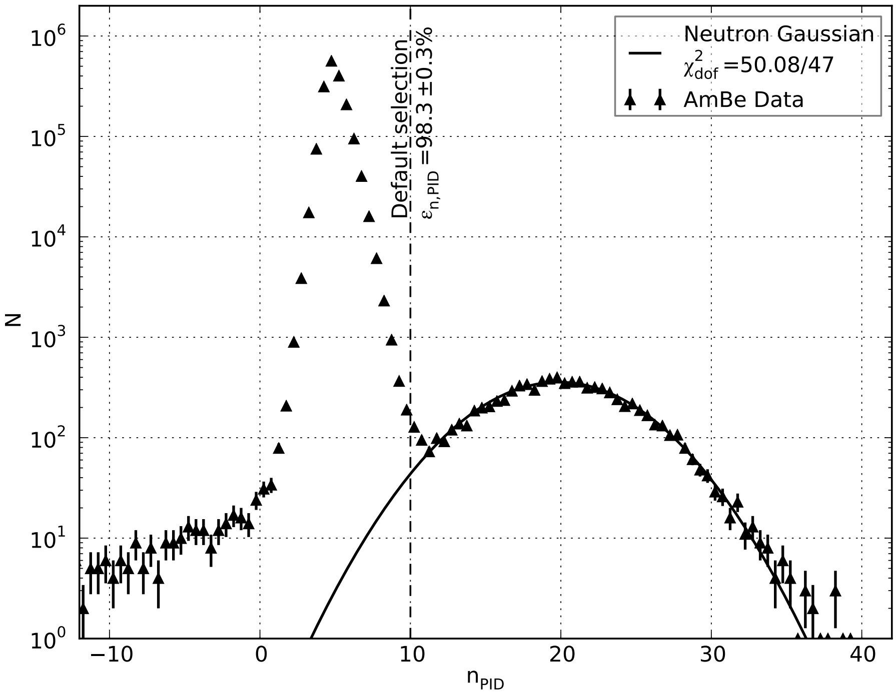

Figure 4.12: Neutron PID parameter for AmBe data with Gaussian fit.

Integration Range Optimisation

One free parameter in the reconstruction chain is the integration range of the waveforms. This value has optimised for EM-ZnS separation. The figure of merit between the two distributions has been scanned as a function of integration range - see [45]. A further condition is imposed that the range be a factor of the period of the electronic pick-up, 8 samples [53], so that it may cancel-out. As a result of these investigations, a default integration range of 40 samples is used.

Remarks for Phase 1

As discussed in the previous chapter, the trigger of the prototype only read-out large amplitude neutrons. For phase 1 of SoLid, since it is more challenging to separate the EM and neutron populations at lower amplitudes, more sophisticated algorithms are being investigated (such as deep learning). For the trigger of phase 1, such algorithms are too complex to be implemented in the trigger electronics (as is this method presented). Simpler neutron ID algorithms, which can be implemented at the trigger level but with reduced purity, are investigated in chapter 7.

4.3.2 IBD Prompt Candidates

It is important to ensure that the reconstruction of events outlined above is appropriate for prompt IBD events, i.e. the positron and annihilation gammas (treated as a single event). Recall from chapter 3 that the light detected by the positron is due to scintillation in the PVT, allowing for an accurate estimation of the initial positron energy. Conversely, light detected from the annihilation gammas arises from Compton scattering by electrons, which typically occur far (i.e. many cubes) from the positron annihilation point. These gammas are always at a fixed energy of 512 keV, although any energy detected from Compon scattering will vary significantly due to quenching [4].

The strategy at SoLid is to reconstruct the positron energy separately from the annihilation gammas. To ensure reconstruction is optimal, simulation is used to compare the initial positron energy with that reconstructed. The simulation includes all detector response effects, such as SiPM cross talk, light yield, attenuation (for more details, see [71]). Two estimators of the positron energy are presented to optimise the reconstruction:

- **Maximum cube energy**: this corresponds to combining the largest energy signals found on horizontal fibres with the equivalent for vertical fibres. This estimator fulfils the requirement of only using the energy local to the positron, however has the disadvantage that the 'quasi' 3D readout (i.e. channels readout the projections of cubes along fibres, not individual cubes) can introduce degeneracies in finding the correct fibre combination. Two events demonstrating this difficultly are shown in figure 4.13.
- **Total energy deposited**: this is equivalent to the event energy outlined above. This estimator is advantageous in the case of SM1, where the detection efficiency of annihilation gammas is negligible, and therefore all energy detected most likely corresponds to the position. In future iterations of the experiment, this will have to be adjusted to only use information local to the positron, for example, by summing signals spatially near the largest energy deposit on each detector array.

The correlation between initial and reconstructed energy for both estimators is shown in figure 4.14. The former estimator has many outliers compared to the latter, and these events are cases where the positron deposited energy in multiple cubes, giving an underestimated energy value. The width (along the diagonal) of the two distributions is similar, and

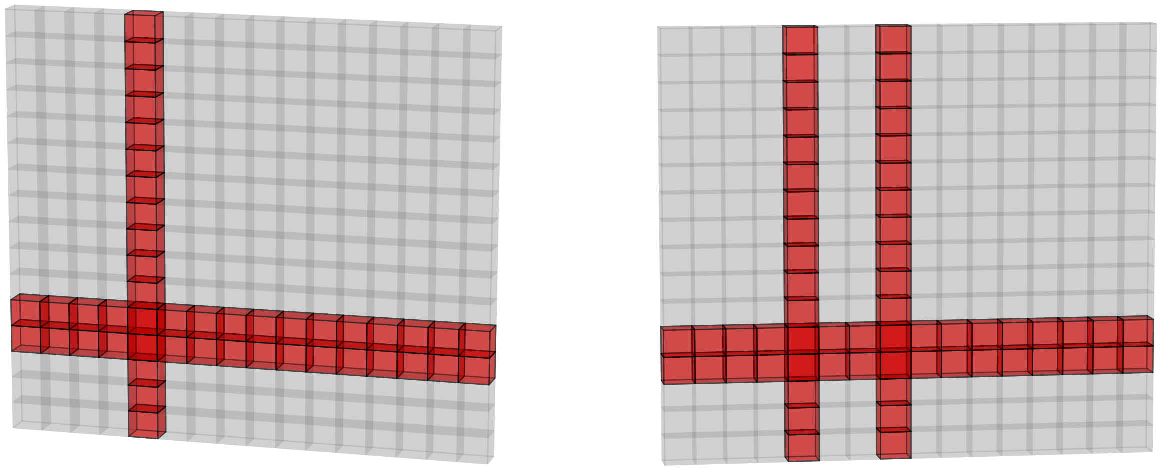

Figure 4.13: Example fake events where the read-out of cubes using fibres can lead to degeneracies when trying to localise signals to specific cubes. Cubes connected to triggered channels are highlighted in red. Left: case where two cubes in the same plane and same row led to a trigger - assigning the energy value of each cube is difficult. Right: case where energy was detected from multiple cubes in the plane, leading to degeneracy in assigning the cube position and energy.

corresponds to the expected energy resolution of the prototype detector ($\sim 20\%$ at 1 MeV). To avoid these outliers, the estimator using the total energy deposited in the detector is retained for prompt candidates.

There are many backgrounds that can contaminate IBD prompt events, including muons, fast neutron proton recoils and reactor induced radiation. Methods to specifically separate muons from IBD prompts are defined in the next section. Further selections to purify the IBD prompt selection are outlined in the next chapter, where the distributions of the event features are also shown.

4.3.3 Muon Candidates

Efficient muon identification is critical, especially for the inverse beta decay analysis. As well as contaminating prompt candidates, muons can produce neutrons via muon spallation, which can increase the backgrounds of the IBD signal (background studies suggest that a muon tagging efficiency above 90% is sufficient for the prototype to ensure this background contribution is relatively small). Muon identification should also be pure, to ensure prompt IBD events are not tagged as muon events, therefore lowering the overall IBD efficiency. This particular separation is now studied in detail using simulations of both signals. Details of

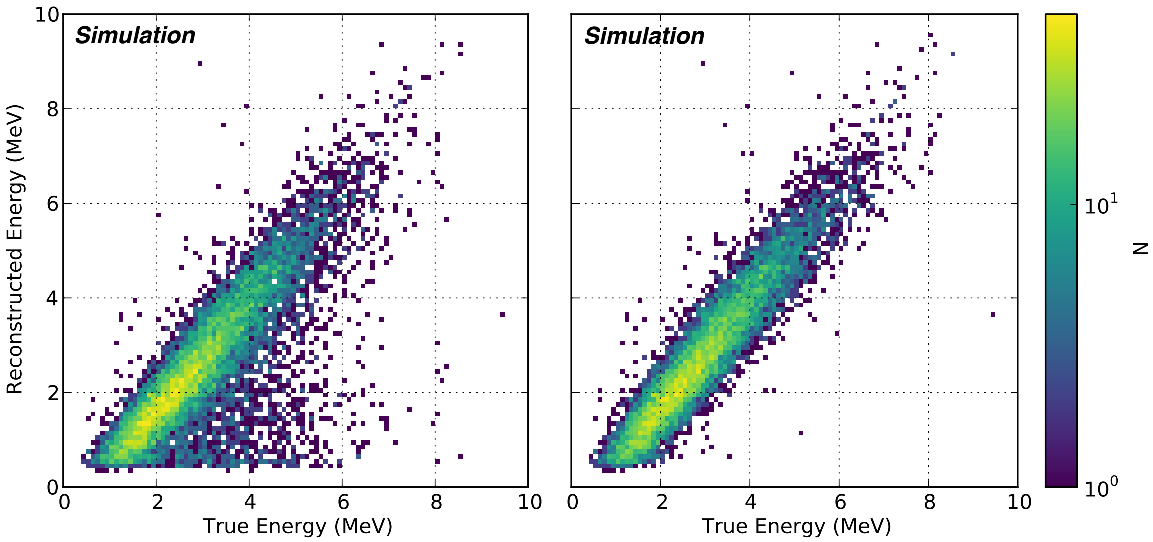

Figure 4.14: Comparison of different energy estimators for IBD prompt reconstruction. This uses simulations of positrons forming IBD candidates, reconstructed using the same algorithms as used for data. Left: max cube, right: total energy deposited in the detector.

the simulations themselves are provided in later chapters.

IBD Prompt and Muon Separation

Compared to prompt IBD events, muons deposit larger amounts of energy and across many more channels. The majority of them cross through the detector, and these can be easily removed. However, some muons may only pass an edge or corner of the detector, and are more likely to be misidentified as prompt events. Consider the 2D distribution of energy vs. volume, for simulated IBD prompt and cosmic muons events, as shown in figure 4.15. The one-dimensional projections, including the cumulative distributions, are shown in figure 4.16. The vast majority of IBD prompt events are found at relatively low values of energy and volume, whereas only a small proportion of muons are found in these regions.

Both variables are explored as a discriminating parameter separately. A useful tool to judge the separation of two populations is the 'receiver operating characteristic' (ROC) curve. This shows the efficiency vs. fake positive rate for a particular selection, often as a scan over a discriminating parameter. These are shown in figure 4.17. Of the two parameters, energy is clearly the more powerful discriminator, as demonstrated by the increased area under the energy curve. Further separation can be gained by considering the spatial distribution of these low energy muon events, for example, by also vetoing events if they intercept the edge

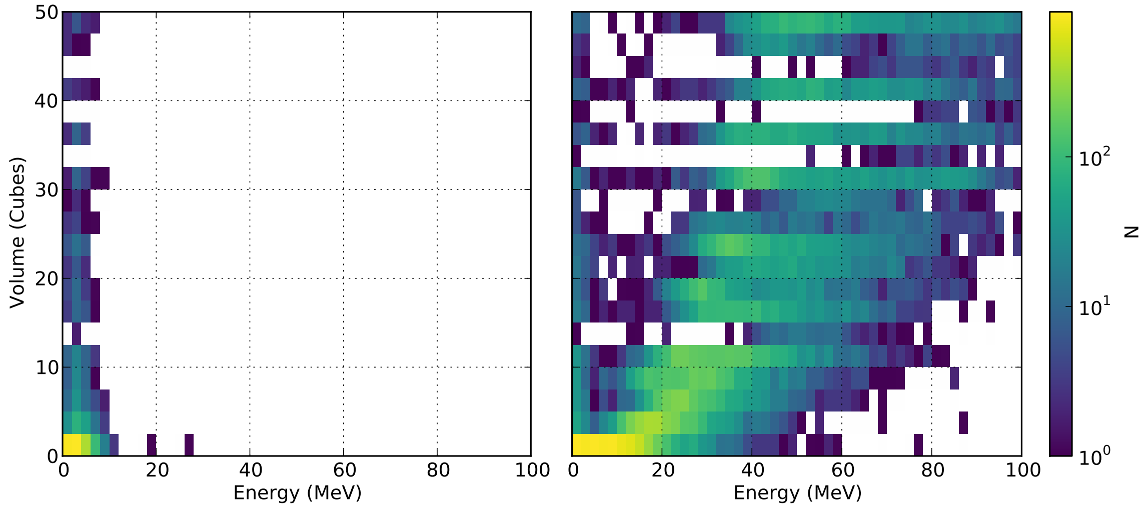

Figure 4.15: Correlation between volume and energy variables for simulated IBD prompt events (left) and cosmic muon events (right). Note that volume is a discrete variable. These distributions are made prior to any additional IBD selection criteria.

cubes of the detector. It is found that this extra condition can reduce the background from muons by approximately another factor of 2 for the equivalent IBD prompt efficiency.

In practice, the default selection criteria for a muon is $E_{Event} > 8\,\text{MeV}$, and without any spatial considerations (this is considered for future iterations of event reconstruction). Using figure 4.17, it can be seen this has a muon selection efficiency of 94%. Lowering this selection would reduce contamination of IBD prompt events further, however, the IBD efficiency would also begin to decrease. This would also have implications for the sterile oscillation search, which uses the shape over the full range of the IBD prompt energy distribution.

Chapter 5 outlines further selection criteria to find muons that can be tracked through the detector to high accuracy.

Data-Simulation Comparison

To validate the prototype simulation, and as a teaser for the next chapter, the EM event energy distribution for reactor on and off data is shown in figure 4.18. Good agreement is observed between simulation and data for both rate and shape at the high end of the energy spectrum, suggesting these events are indeed purely muons. The IBD contour is scaled to

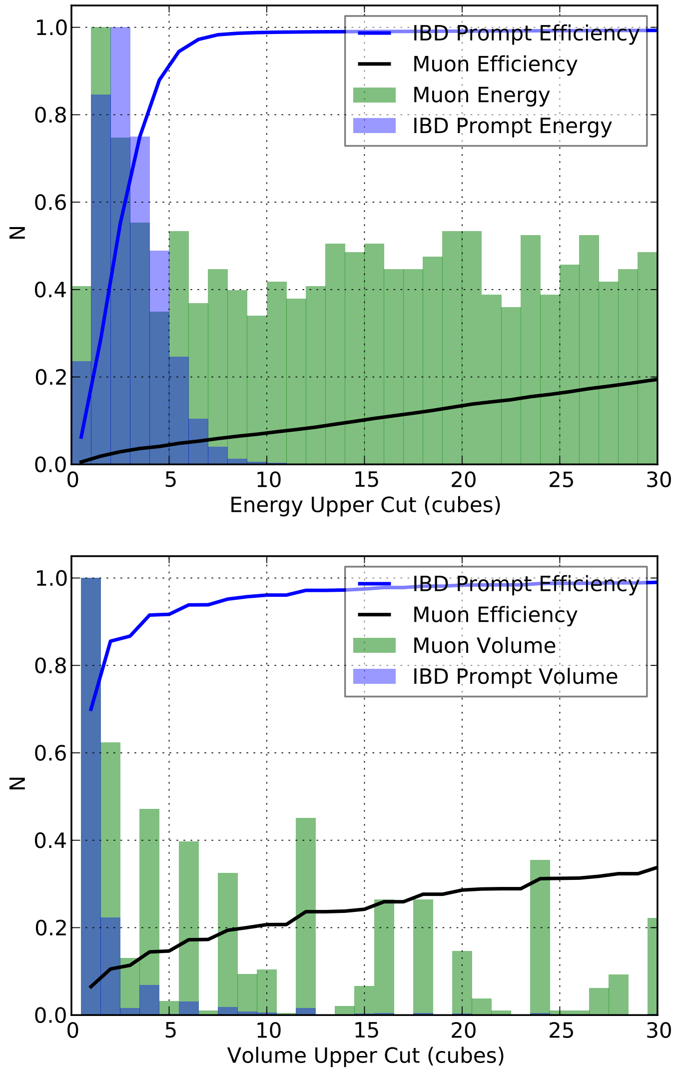

Figure 4.16: Distributions of event volume and energy (in MeV) variables for simulated IBD prompt and cosmic muon events.

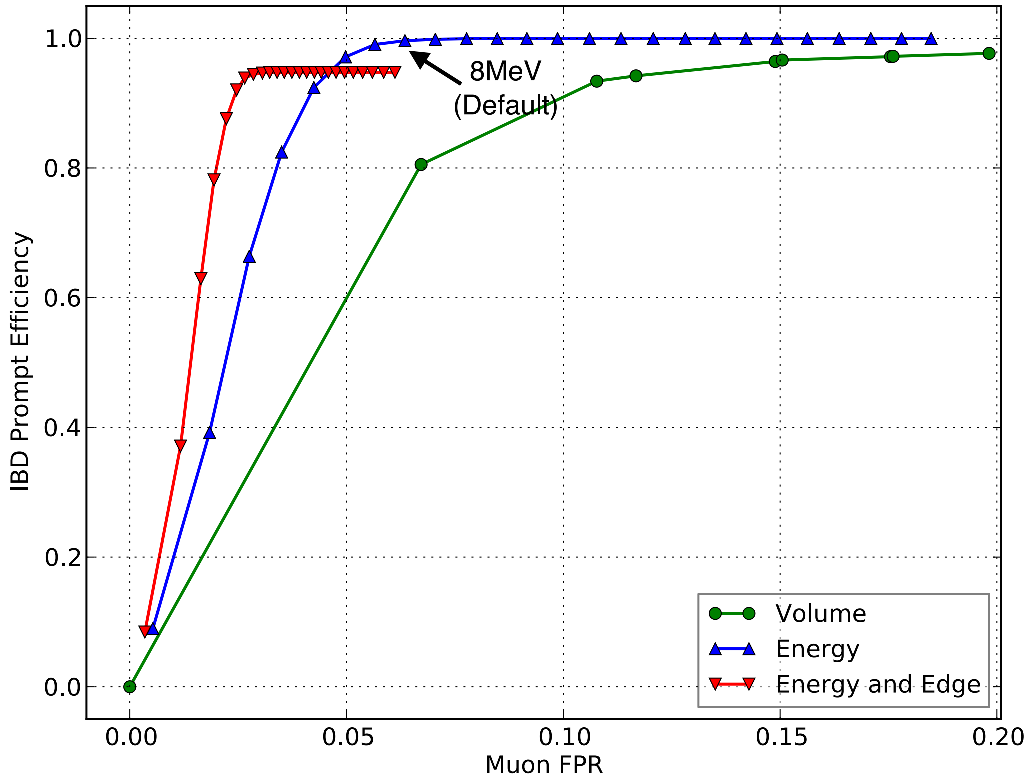

Figure 4.17: ROC curves for three different discriminators used to separate prompt IBD events from cosmic muons.

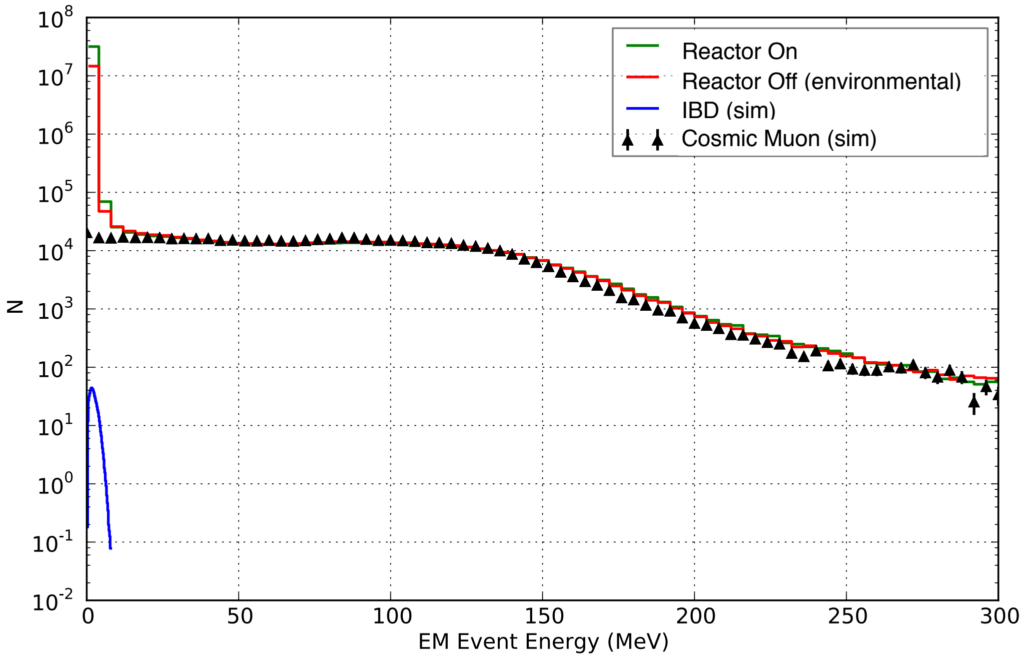

Figure 4.18: Energy spectrum of electromagnetic events for the reactor on and reactor off periods (scaled for equal live time). The simulation prediction for IBD prompt and muon events is overlaid.

the expectation (outlined in the next chapter), and prior to any IBD analysis selections, the signal to background ratio is approximately one to one million. This demonstrates the true magnitude of the challenge of this kind of experiment. Of course, this histogram is projected over all of the EM event feature phase space, including regions which are known not to include any signal at all. This is refined in the next chapter. Further data-simulation comparisons can be found in [71].

4.4 Electromagnetic Calibration Application

Energy calibrations are applied at the peak finding and event finding stages of reconstruction. At peak finding, per channel corrections, such as gain equalisation, can be applied; at event finding, per cube corrections can be applied, using the 3D position information. This section briefly describes the application of these calibrations. The methods themselves are outlined more completely in chapter 5, including methods used to find the absolute energy scale of the detector and calibration values.

4.4.1 Channel Calibration

Each silicon photo-multiplier has a unique breakdown-voltage. In order to equalise the energy response of the detector, the gain of each channel has to be calibrated individually. Online gain calibration was performed during the commissioning of the large scale prototype, and is accurate to a relative spread of approximately 20% across channels. This spread can be reduced to $< 3\%$ by introducing a calibration term for each channel α_i, which linearly scales the integral response of the ith channel. Another term β_i is introduced to linearly equalise the I/A response (i.e. a shaping parameter), to account for the variation in waveform shapes between channels. These are applied to all peaks as follows:

$$I'_{Peak} = \alpha_i \times I_{Peak} \tag{4.1}$$

$$A'_{Peak} = \alpha_i * \beta_i \times A_{Peak} \tag{4.2}$$

where $'$ denotes the calibrated values.

4.4.2 Cube Calibration - Attenuation

The light yield of each cube depends on the distance between the cube and the end of the optical fibre. Light is lost due to attenuation, and it is found that the light yield can vary by approximately 20% when comparing one of the fibre with the other. Using the three dimensional information of an event, light loss due to attenuation can be corrected, providing a uniform energy response along the fibre (the energy resolution will still depend on the position along the fibre).

The relative light loss F_{Loss} is modelled as the sum of two exponential terms: the first represents light that travels directly from the event to the silicon photomultiplier, and the second term represents light from the reflection on the mirror. This is written as:

$$F_{Loss}(x) = \frac{1}{2}e^{-x/a} + \frac{1}{2}re^{-(2X-x)/a} \tag{4.3}$$

where x is the position along the fibre of the event, and X is the length of the fibre. r and a are constants whose values and meaning are described in chapter 5. Using this equation, attenuation is corrected by solving for $F_{Loss}(0)$ for each hit forming the event. That is to say, energies are corrected such that there is no loss of light due to attenuation. This is applied to all IBD prompt candidate events, since these are localised to a small number of touching cubes, where r can be accurately found (this is not always the case for larger events, such as muons).

4.5 Implementation in Saffron

All the above reconstruction is implemented in a new dedicated framework for SoLid reconstruction and analysis, called Saffron (**S**oLid **A**nalysis **Fr**amework[1]). Written in C++, Saffron is used for large data processing jobs. The successor of this software project, Saffron2, is currently being written for phase 1 data taking, and will be used in multiple instances including online triggering and filtering, offline reconstruction, and other calibration and analysis tasks. This section briefly outlines the implementation.

4.5.1 Processing Cycle

Similarly to other reconstruction packages such as Gaudi [70], Saffron involves processing data in a large reconstruction cycle. Raw data is loaded into memory for a fixed buffer that corresponds to a fixed interval in data taking time. The interval is made very large relative to the time-scales of the data, to allow users to find correlations at those scales without concern for memory management. Saffron uses a time interval of 100ms - large enough to avoid to a significant number of clipped events at the boundaries of the buffer.

4.5.2 Saffron Algorithms

The above reconstruction chain is split into several reconstruction algorithms. These are C++ classes that inherit the same structure, and can be enabled or disabled as the user wishes.

4.5.3 Data Streams

Saffron has been used to generate multiple sets of offline reconstructed data via various output data streams for different analyses. These data streams differ depending on the output save options. For example, for the IBD analysis, the output of Saffron reconstruction can be limited to short time-windows around neutrons only. This gives a large data reduction factor

[1]The author admits this naming decision was made during a finale of Masterchef.

$\mathcal{O}(1000)$. For muon studies, there is a a smaller but still significant reduction factor $\mathcal{O}(100)$, due to the higher rate of muons compared to neutrons. Smaller files allow for quicker analysis, and since the output formats are ROOT based tuples, this is more user-friendly than raw waveform data, allowing many non-experts of the software to study the data.

4.5.4 Event Display

Saffron also includes an event display, which uses the python library MatPlotLib [44]. The user defines a specific time window, of variable width, to view all raw and reconstructed data. As well as displaying raw waveforms, the position of the triggered SiPMs is also shown, alongside a three-dimensional view of the detector that shows the extrapolated event and topology. A time-line displaying the time-stamps of events is also included. An example event display is shown in figure 4.19. Examples of IBD event displays are shown in the next chapter.

4.6 Conclusions

Reconstruction of SoLid raw data has been performed at the time-scale of the timing resolution of the detector. This applies to interactions inside the detector that give rise to simultaneous signals. Raw waveforms are cleaned using several techniques, and SiPM pulses are found by peak finding. These peaks are time clustered to form events, and a rich set of event features are extracted, including many topological features. Calibrations have also been applied at various stages of the reconstruction to optimise energy reconstruction, correcting for SiPM variations, and attenuation loss in the optical fibres. Particle identification has been performed to categorise these events into neutrons, muons and other EM events (such as IBD prompts). Effective categorisation is critical for the IBD analysis. The energy reconstruction of prompt IBD events has been tested with multiple estimators, and it has been demonstrated that sufficiently accurate energy reconstruction is possible using this technology.

All of these stages have been implemented in a new dedicated software package for SoLid reconstruction and analysis called Saffron. Saffron has been used to generate several reconstruction data streams as input to many other studies, and as such, its development is now an intrinsic part the SoLid analysis effort. This has led to the creation of a new project, Saffron2, which is a significant re-factor of is predecessor. Saffron2 will be used in many data taking stages during future phases of the experiment, including online triggering and

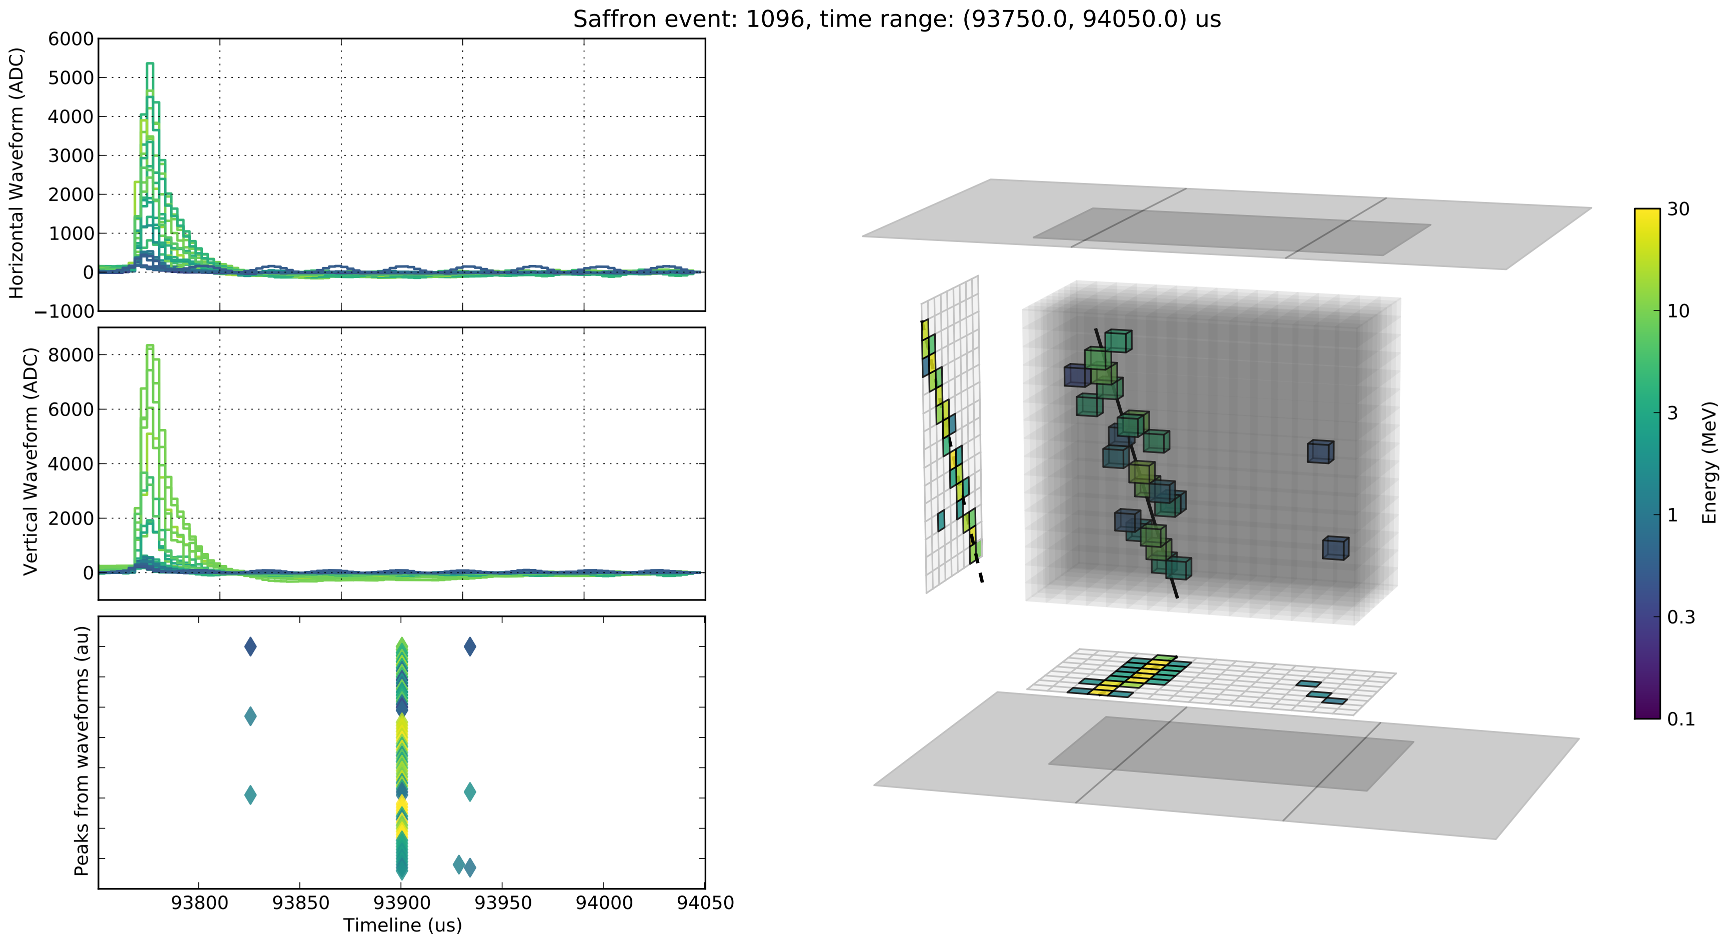

Figure 4.19: Example event display of a muon candidate, using Saffron1 reconstructed data as input. The program is written in python, and uses the Matplotlib library [44] for graphics. EM events use the colour-scale that shows the energy of the hit forming the event. The waveforms of events, split by fibre direction, are shown in the top left (with arbitrary phase, so they all line up). For timing information, the time-line in the bottom left can be used. The data plotted is for a pre-specified time window, that can be varied. The muon veto panels are optionally displayed. The event display was developed by the author.

filtering.

Chapter 5

Muon Energy Calibration

5.1 Introduction

Muon events are abundant at SoLid. They dominate the upper region of the EM event energy spectrum, and can be effectively identified from other kinds of events by their large amplitude signals and high channel multiplicity. As well as being a source of background in the neutrino analysis, they can be used elsewhere, including energy calibrations. These calibrations are required at two instances: per SiPM channel, and per cube.

This chapter outlines a new calibration procedure for EM signals, developed for the prototype SoLid module, using tracked muons. It is structured as follows: the introduction continues by describing the requirements for energy calibrations. Section 5.2 describes methods to track muons through the detector, and section 5.3 outlines the selection criteria used to refine the default muon selection for calibration purposes; specifically, in order to find the distribution of energy deposited per unit path length for each cube. Section 5.4 then shows how these distributions can be used to equalise the detector SiPM channels, and to correct for attenuation losses due to the optical fibres. The remaining variations at the cube level are explored, and the absolute energy scale of the detector is found from comparisons with simulation. Finally, the method is validated using simulation in section 5.5. Unless otherwise stated, the beginning of the reactor off dataset is used.

5.1.1 Requirements

Energy calibration procedures at SoLid are required to perform the following:

- **Channel equalisation (i.e SiPM and fibres)**: the online channel equalisation performed during commissioning, required since each SiPM has a unique break-down voltage, are accurate to $\sim 20\%$. This can be improved offline.

- **Cube equalisation (i.e attenuation and fibre-cube coupling)**: light is lost due to attenuation in the optical fibres. This can be corrected using the 3D information of events. Additionally, significant variations have been previously measured in the coupling of the fibres to the cubes.

- **Light yield measurement**: methods are required to convert the amplitudes of SiPM EM signals into absolute energy measurements.

The maximum energy resolution possible for the prototype (due to light yield considerations alone) is limited to $\sim 20\%$ at 1 MeV [53]. Equalisations should be accurate enough to not significantly worsen the energy resolution further due to smearing across cubes/channels, although need not be so accurate to the sub-percent level given this limitation. In practice, equalisations accurate to the 5% level are considered acceptable for the prototype. Further, the only component of energy calibration expected to be environmentally dependent are the SiPM calibrations. Since their break-down voltages are correlated with temperature, these calibrations may need to be repeated periodically, depending on the speed of temperature changes. The other components are expected to be stable in time, although are still monitored regularly.

5.2 Muon Tracking

5.2.1 Track Finding

It was shown in chapter 4 that muons can be effectively tagged by selecting events whose total energy deposited in the detector is above 8 MeV. It was also shown that energy is highly correlated with the volume parameter, which can be found from topological information alone, thus allowing muons to be identified prior to the application of energy calibrations. An example of such a muon is shown in the event display in figure 5.1. As can be seen, muons form clusters of triggered channels (correlated in time and space) on the arrays of horizontal and vertical fibres.

For tracking purposes, the path of the muon is approximated as a straight line through the

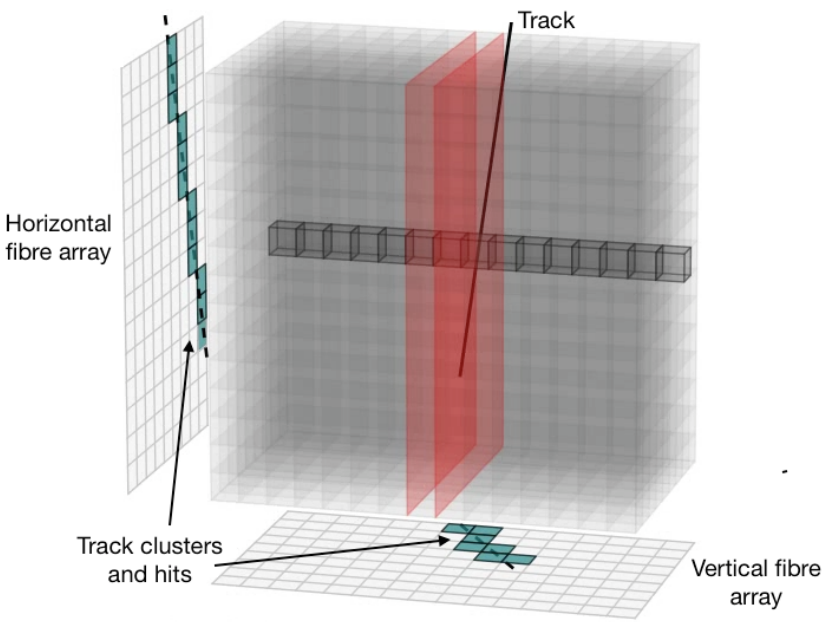

Figure 5.1: Example muon event used for calibration. The clusters of triggered channels are highlighted on the two SiPM arrays as gray, and the results of the linear regression to fit a straight line to these shapes are shown as the dashed lines. The full line is the 3D extrapolation, which can be seen to pass through just one cube (marked by the red plane) from the highlighted row of cubes.

detector. This is a valid approximation, since the deviations from a straight line due to multiple scattering inside the detector are very small for PVT [1]. In practice, the input used for tracking are the shapes of the clusters on the two fibre arrays. Straight lines are fitted to these two shapes independently. Although this method does not necessarily use the full 3D output of the detector, working with the projections greatly simplifies the tracking procedure, and as shown below, the performance of the algorithm meets the requirements.

A minimum of two points is required to fit a straight line using linear regression. Using the features extracted for each event (outlined in chapter 4), this corresponds to the requirement:

$$\mathrm{Min}(\mathrm{size}_{xz}, \mathrm{size}_{yz}) = \mathrm{size}_{Min} \geq 2 \tag{5.1}$$

where size_{Min} denotes the minimum cluster size of the two SiPM clusters (recall that size_{ij} is the number of triggered channels in a particular optical fibre direction ij). Of the initial selection of muons, around 97% fulfill this requirement, allowing them to be tracked. The accuracy of the tracking is discussed later in section 5.2.2. The distribution of size_{xz} v.s.

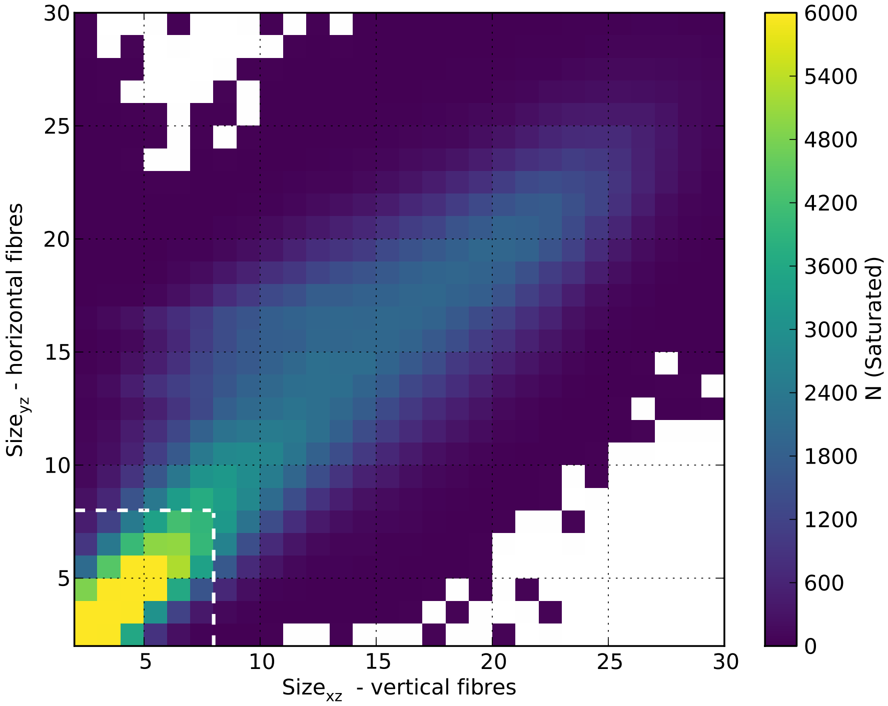

Figure 5.2: The distribution of size$_{xz}$ v.s. size$_{yz}$ for muon events. The colour bar is saturated to also show the structure at higher values of size$_{yz}$. The white lines show the cuts applied to select tracks with many hits, which can be reconstructed with greater accuracy.

size$_{yz}$, for events from data that fulfil this criteria, is shown in figure 5.2. There is a strong correlation, with a bias towards larger size$_{yz}$ (i.e. larger clusters formed of horizontal fibres). This is expected, since on average the direction of muons are closer to vertical than horizontal, resulting in smaller clusters on the vertical fibre array and larger clusters on the horizontal fibre array.

There are a small number of cases ($\sim 5\%$) where tracks are contained within a single detector plane only - see figure 5.3 for examples. These cases are troublesome, since there are multiple possible incident angles of the track. As such, tracking the muon through specific cubes is unreliable, and so these cases are removed from calibration studies. In future iterations of the method, timing information can be used (based on time-of-flight measurements) to resolve the incident angle.

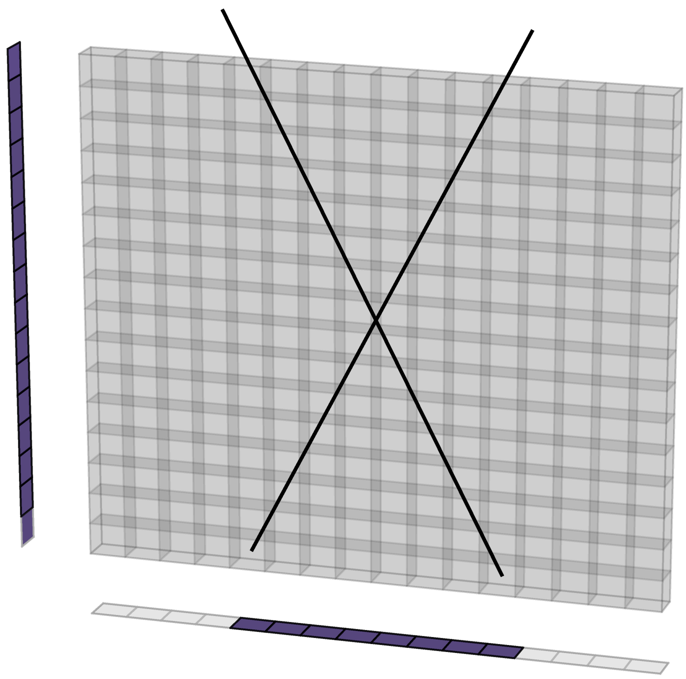

Figure 5.3: Example muon event that passes through a single plane only between entering and exiting the detector. Since there are two possible angles for the muon (shown in black lines), the cubes the muon pass through cannot be determined with certainty using tracking information alone, therefore these cases are removed.

5.2.2 Tracking Resolution

As with usual straight line fitting, the accuracy of the fit increases with the number of the data points available. For muon calibrations, it is important to be able to accurately reconstruct the pathlengh of muons through specific cubes: dx. The accuracy of the method to find dx is shown in the top of figure 5.4 using simulated tracks. This shows the residual distribution: Δdx, which is the difference between the fitted dx and original dx from the simulation. This has been drawn as a function of size$_{Min}$.

The expected trend, that the reconstruction accuracy increases with the number of data points, is observed. To quantify this effect, consider the RMS of the projection of this distribution, projected between each value of size$_{Min}$ and infinity. This is shown in the bottom half of the figure. The RMS is seen to drop from 0.7cm when using all tracks (i.e. size$_{Min}$ $\geq$ 2), and plateaus around size$_{Min}$ $\geq$ 8 with an RMS of 0.5 cm. This increased accuracy comes at the expense of rejecting near 50% of entries, as shown in the acumulative distribution also plotted. This plateau region is used for calibrations, for increased accuracy and fewer outliers.

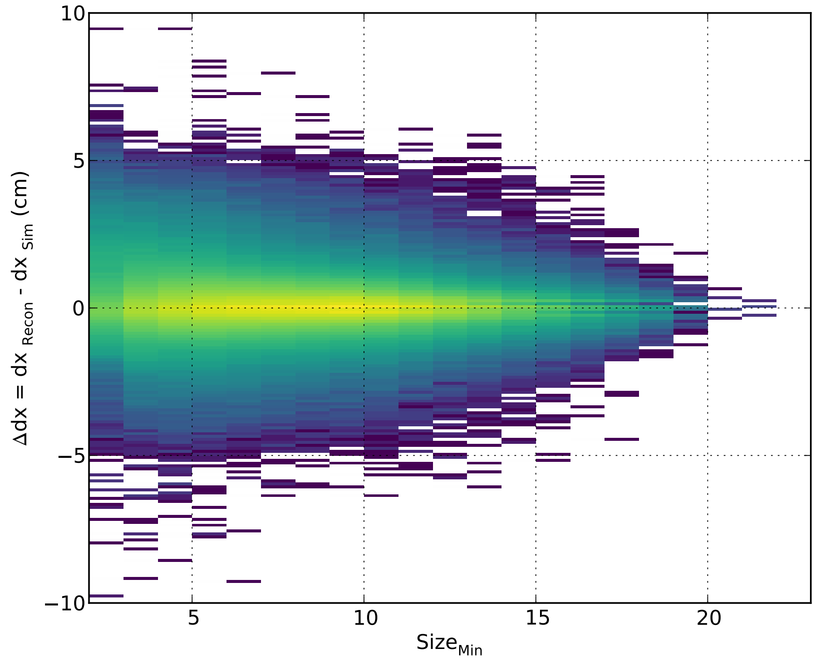

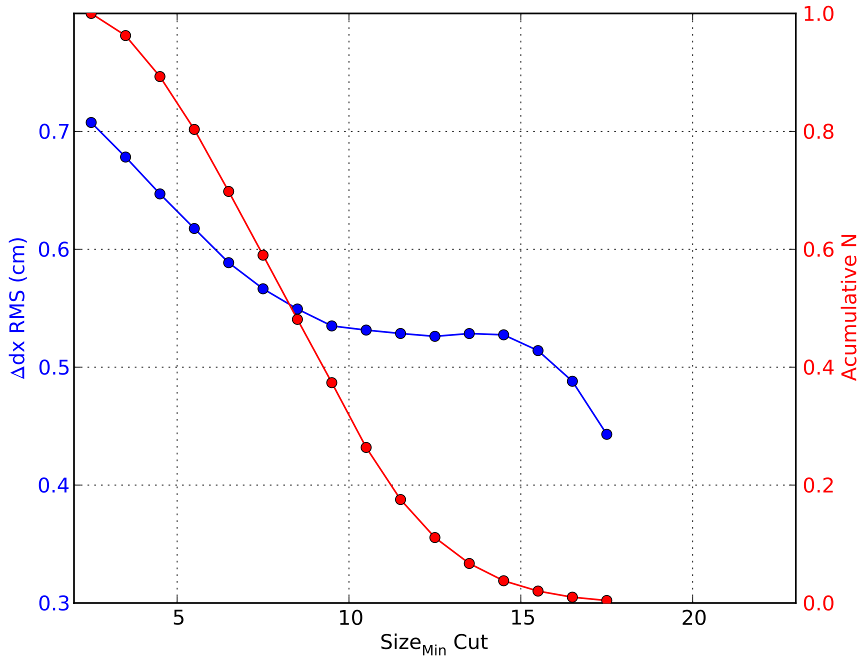

Figure 5.4: Top: residual distribution of simulated tracks between generated and reconstructed dx. Bottom: the RMS of the residuals when placing lower limits on size$_{Min}$, and the corresponding acumulative (i.e. 1-cumulative) distribution to show the selection efficiency.

5.3 Finding dE/dx

The energy deposited per unit path length dE/dx is an important variable for energy calibration, used at several other experiments [56]. As well as being able to accurately reconstruct muon tracks, additional selection criteria are required to find energy deposits that can be localised to specific cubes. These are now outlined.

5.3.1 Cube Containment

An example muon event, where part of the energy deposited by the track can be localised to a single cube, is shown in the event display in figure 5.1. Consider the row of cubes highlighted in grey. In the example shown, the track can be seen to pass through a single cube only in the row. In this case, we can say with certainty that the light output from the fibre in this row corresponds to this cube only. Such certainty is not always guaranteed - in the vertical direction, the muon traverses many cubes along the same optical fibre.

Geometrical selection criteria are used to select these cases. For each energy deposition on an optical fibre, the muon is propagated through the corresponding row of cubes, and selected if the number of cubes intercepted is equal to one. The value of dE/dx for the cube is then taken as the ratio of the energy deposited on the fibre, and the path-length through that cube. The fraction of channel hits that fulfil this condition is around 40%.

5.3.2 dx Bin

The dE/dx distribution, for muons travelling through scintillator at a *fixed pathlength*, can been parametrised as a convolution of a Landau and Gaussian distribution. For a detector with perfect energy resolution, dE/dx follows a Landau distribution [1]. The Gaussian convolution is used to describe the smearing by the energy resolution and tracking ability of the prototype detector. A bin of dx is chosen to fix the pathlength. Whereas it is possible to use all values of dx, the resulting shape can be difficult to predict. This is because using a wide range of pathlengths corresponds to a wide range of energy deposits, leading to significant variations in light yield and thus energy resolution. The dx bin itself is chosen to contain a high number of calibration events, whilst not being too wide as to smear the resulting dE/dx distribution. In practice, the bin $1.0 < dx$ (cubes) < 1.2 is used, which contains around 30% of entries. This can be seen in figure 5.5. This figure also shows the correlation between dE and dx, in addition to the projections of dE and dx. A linear relationship is observed

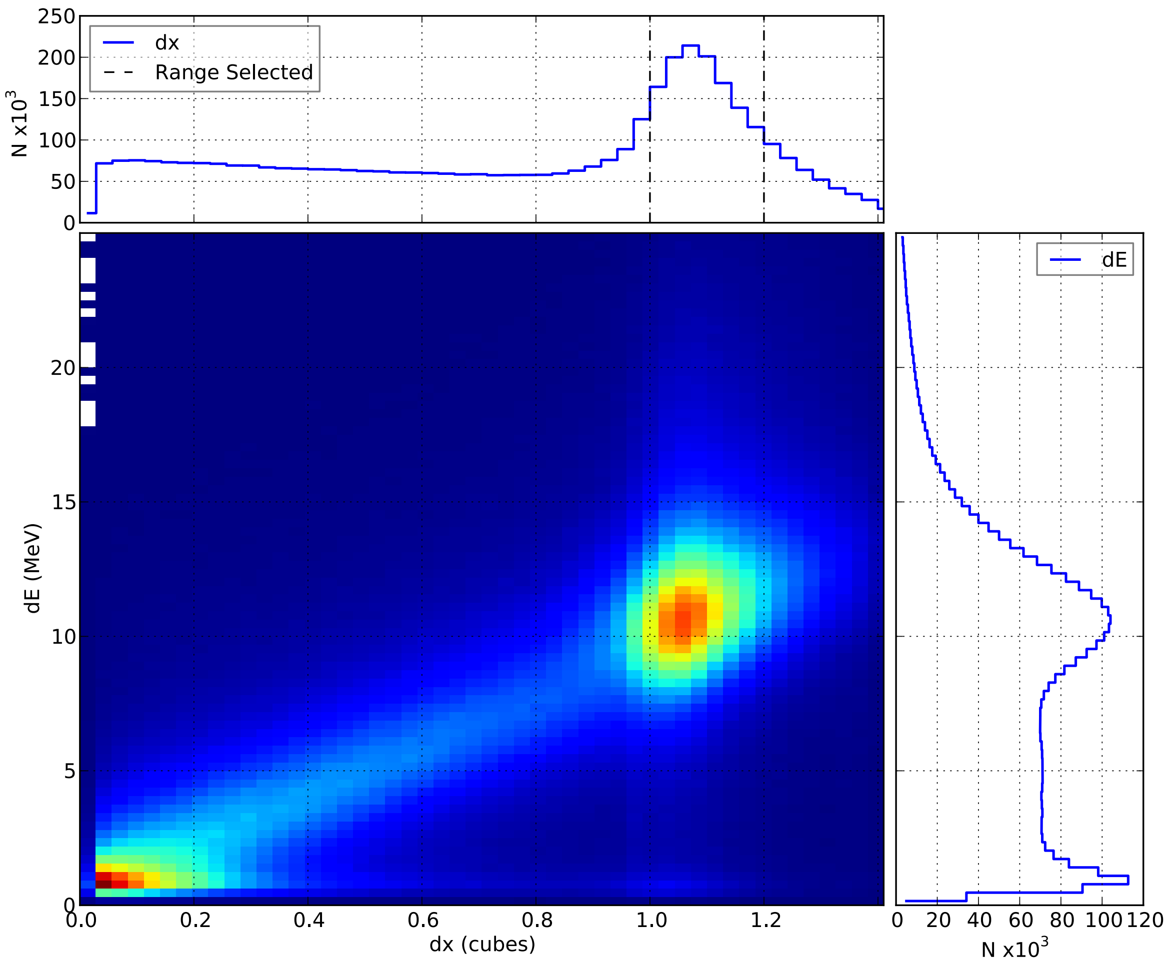

Figure 5.5: Correlation between dE and dx for hits forming muon events, with the projections. The expected linear relationship is observed.

between muon pathlength and the detected energy, as expected for PVT. It is also interesting to note the shape of the dE distribution does not follow a well known shape, demonstrating one of the benefits of using dE/dx instead of just dE.

To estimate the pathlength resolution, the residual distribution of dx after all these selections is shown in figure 5.6 using simulation. The distribution reasonably follows a Gaussian distribution, although there is a small tail at positive values where the reconstructed pathlength has been overestimated. The width of the Gaussian is 0.3cm, or 5% relative to the pathlength bin chosen (5.5cm).

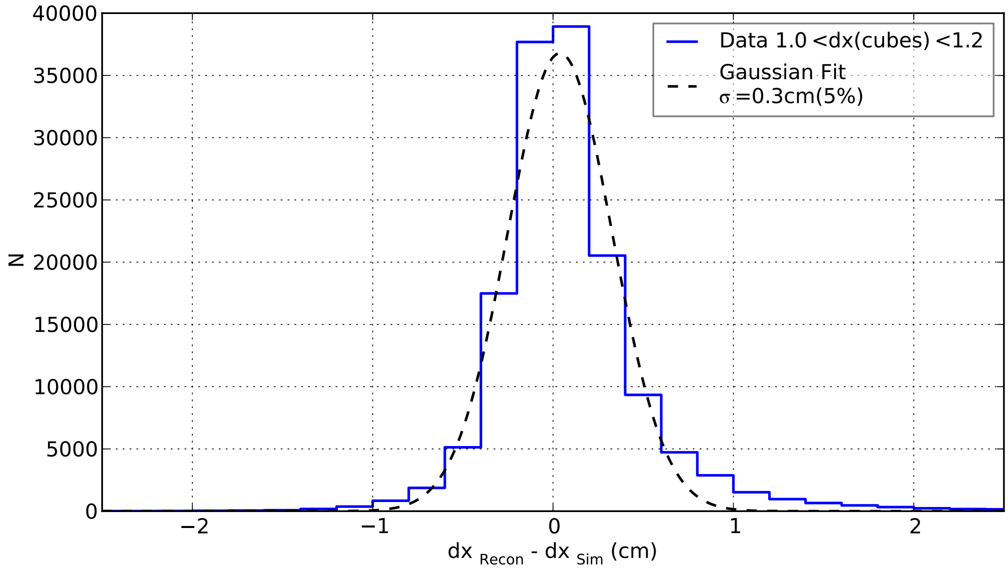

Figure 5.6: The pathlength residual distribution for reconstructed muon paths through individual cubes. This uses simulations of tracks, with all selection criteria outlined in this section applied.

5.3.3 Selection Summary

The above conditions, including the effects on the rates of entries that could be used for calibration, are summarised in table 5.1. Of the initial muon sample, 2.5% of muon hits are used for calibration. Assuming that any calibration method requires at least $\sim$ 1000 hits for each cube for accuracy, this means muon calibrations (in this form) can be performed approximately daily. The resulting dE/dx distribution after these conditions is shown in figure 5.7, for a single cube, after one days worth of data taking (reactor off). A fit of Landau-Gaussian convolution is performed, which gives a reduced chi-square value of 50/32 in the range (1, 5) MeV/cm, showing the peak of the curve has been appropriately parametrised.

A small excess can be observed at low values of dE/dx. Since dx is fixed, these entries must be at low energies. Due to the large environmental background of electromagnetic events at the reactor site, it is possible these low energy entries are tagged accidentally as muon hits (i.e an accidental background, where reactor gamma events have merged into muon events). This is evidenced further in figure 5.8, which shows the correlation between dE/dx and time separation between the track hit and the track itself. The low energy population extends to much larger values of Δt compared to the higher energy entries, supporting the hypothesis

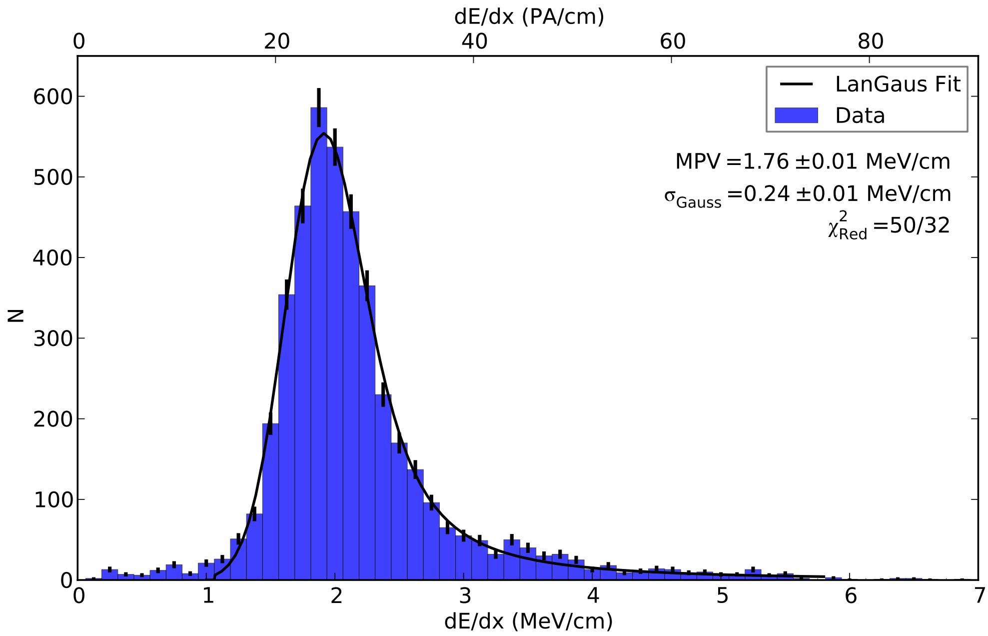

Figure 5.7: Example dE/dx distribution, for a single cube, after one days data taking (reactor off). The top scale uses the number of SiPM pixel avalanches, and the bottom scale shows the energy scale from comparisons with simulation.

that these low energy entries are an accidental contribution. In future iterations of this method, this contribution could be reduced by placing a harsher cut on the time-window used to form events (the current default is 50ns).

5.4 Calibration Application

5.4.1 Channel Equalisation

Using the selection criteria outlined above, the dE/dx distribution for each SiPM channel can be found (i.e. by integrating over all cubes readout by a particular fibre). This is shown in figure 5.9. The channels have been grouped by detector plane. In all planes except the first and last, two distinct groups are observed, depending on the channel number within a plane. These groupings correlate with the type of amplification circuit used for the channels, resulting in two types of waveform shape [53]. These can be seen in figure 5.10, which show examples of average waveforms for channels with different amplifier types - waveforms from ‘Top’ amplifiers are broader, and have a larger undershoot compared to those from ‘Bottom’ amplifiers. This results in a different relationship between integral and amplitude for the two different types of channels. The channels were equalised online according to amplitude

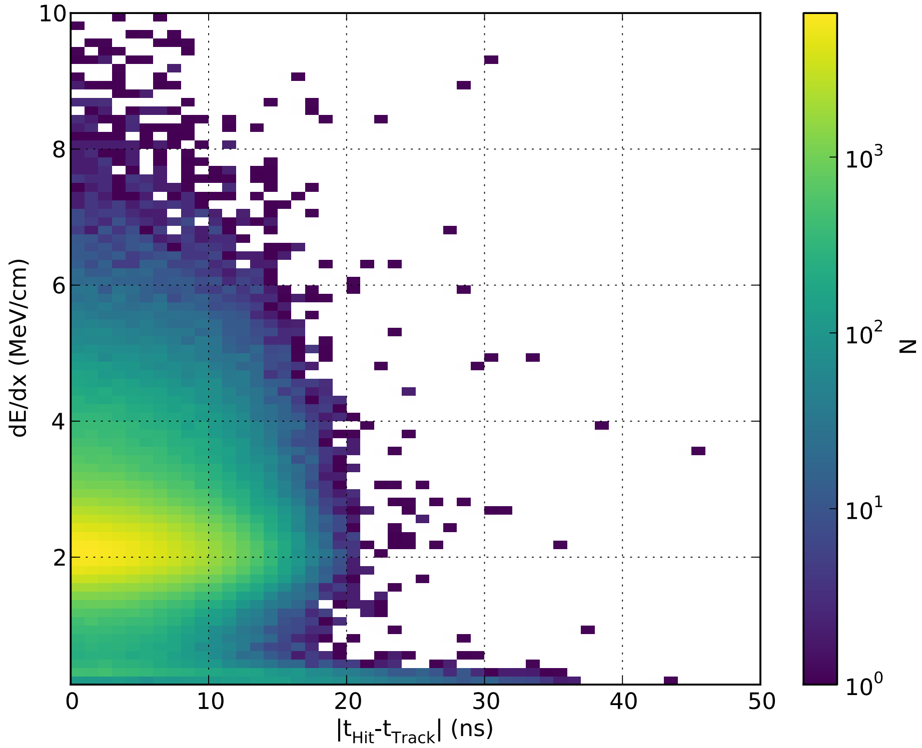

Figure 5.8: Correlation between reconstructed dE/dx and time residual (between the track and calibration hit). The population at low dE/dx extends to high values of time residual, suggesting an accidental component.

Condition	Description	N Muon hits for calibration (1 Day) Channel Mean (x10^3, 2sf)	Cube Mean (x10^3, 2sf)	Selection efficiency
None	All events with energy above 8 MeV	410	51	(100%)
Track fitting	Require size$_{Min} \geq 2$ for track fit	410	51	99%
Multiple planes	Remove cases contained to single plane	370	46	89%
Cube containment	dE can be localised to specific cube	130	17	32%
dx bin	Fixed pathlength	40	5.1	10%
Track resolution	High quality reconstruction size$_{Min} \geq 8$	11	1.3	2.5%

Table 5.1: Summary of the criteria used to find dE/dx with the rate per day and selection efficiencies.

	Channel Group Spread RMS_{MPV}/MPV		
	All Channels	Top Channels	Bottom Channels
Before Equalisation	43%	16%	25%
After Equalisation	3%	2%	2%

Table 5.2: Summary of the spread in MPV values for the dE/dx across channels.

response, although the integral is used in practise for energy measurements (to cancel periodic noise), and it is this effect that gives rise to the patterns observed in figure 5.9. Within the groups of amplifier type, the spread in channel response is $\sim$ 20%, which must be corrected.

The equalisation procedure linearly scales the response of all channels to the channel average. To do this, the Landau most probable value MPV parameter is used, which is a measure of the Landau peak position. The procedure itself:

- Fit the dE/dx distribution to find the MPV for each channel.
- Calculate the average MPV, $\overline{MPV}$, from all channels.
- The linear calibration constant α_i is taken as the ratio of $\overline{MPV}$ to MPV_i.

where i represents the ith channel. For example, if the fitted MPV is twice as high as the average, then the calibration constant is a half, and all energy values from that channel would be reduced by a factor of 0.5×. The spread in MPV values for all channels, grouped also by amplifier, is shown in table 5.2, for both before and after the application of the channel equalisation. This is also shown visually in the pull distributions shown in figure 5.12. Of the 278 active detector channels, 8 failed the fitting procedure. In these cases, the arithmetic mean (which can be found directly from the distribution without fitting) is used for calibration instead of the MPV value, using the linear relationship observed in figure 5.11. The spread in calibration constant (i.e. the RMS) across all channels is 2% post equalisation. This is reduced from 43% prior to the procedure. This can also be seen in the bottom of figure 5.9, which shows the dE/dx vs. channel distribution with channel equalisation enabled.

5.4.2 Attenuation Correction

Light is lost due to attenuation in the optical fibres. This is observed in figure 5.13, which shows examples of dE/dx as a function of cube position along the fibre. An exponential-like

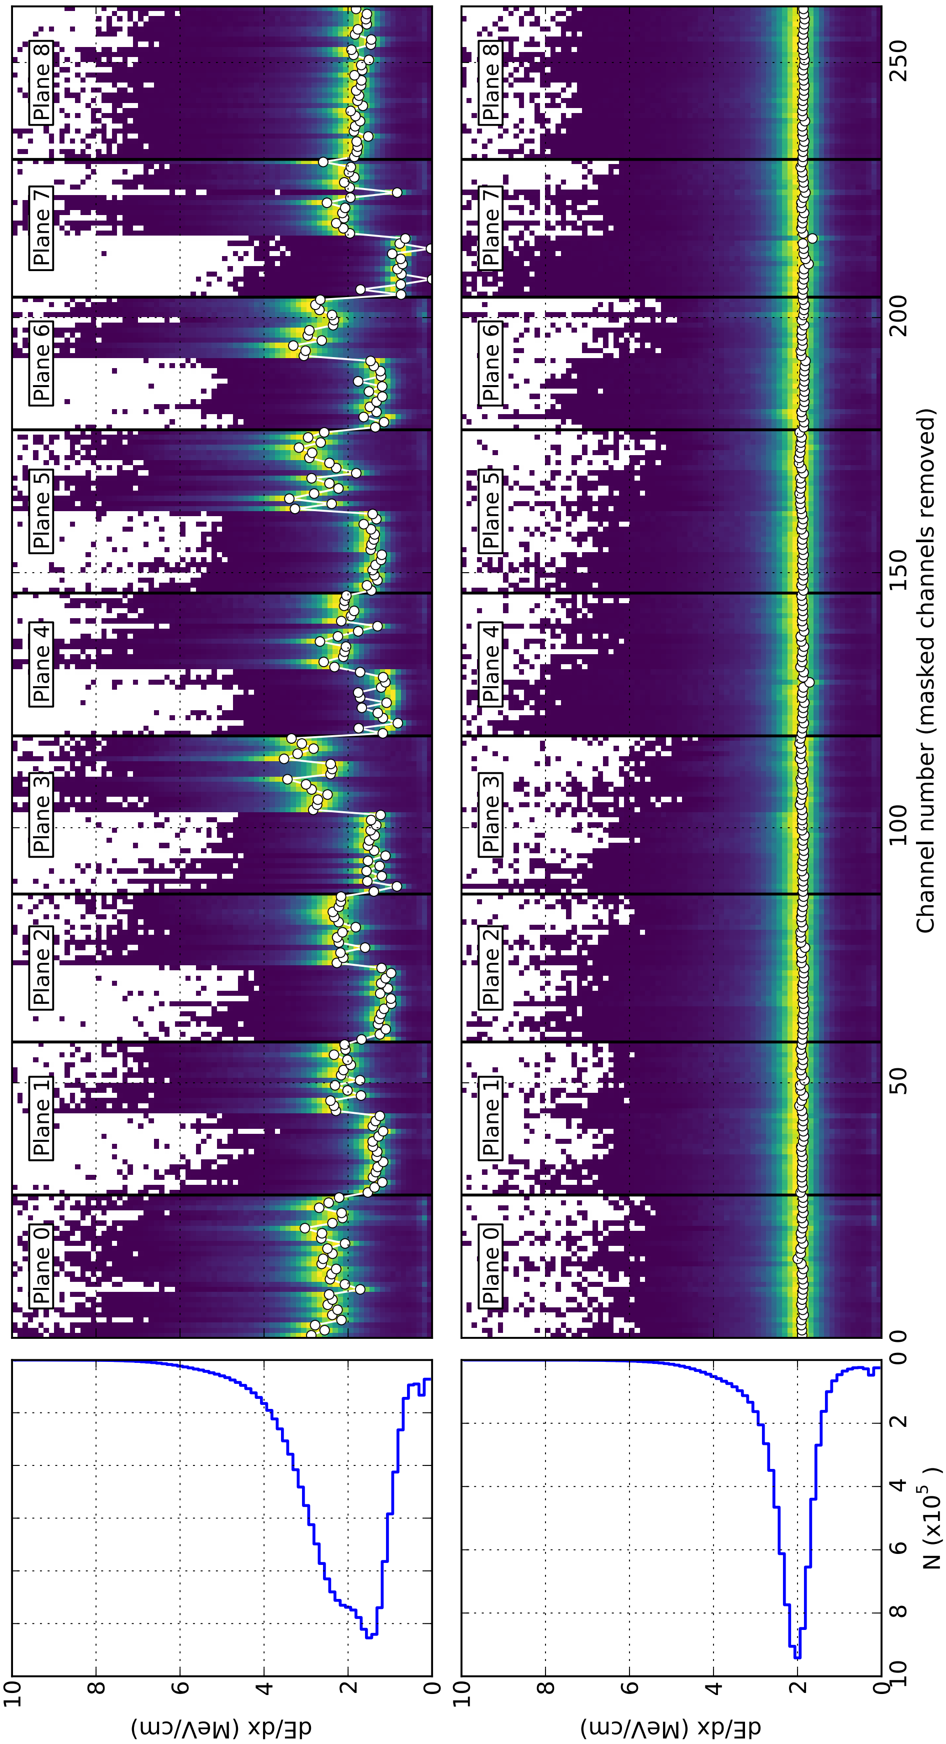

Figure 5.9: dE/dx for each of the prototype SiPM channels, grouped by plane and amplifier type. The white scatter points show the MPV of a fit to each slice. Each slice has been normalised when drawing the plot. Top: before channel equalisation. Bottom: after channel equalisation.

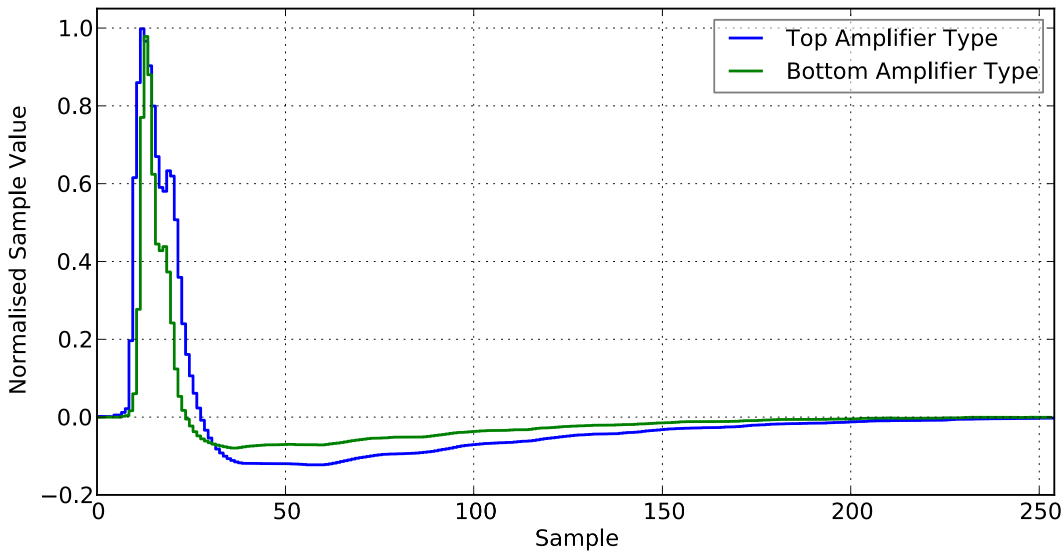

Figure 5.10: Examples of average waveforms from channels connected to different amplifier types (‘Top’ and ‘Bottom’).

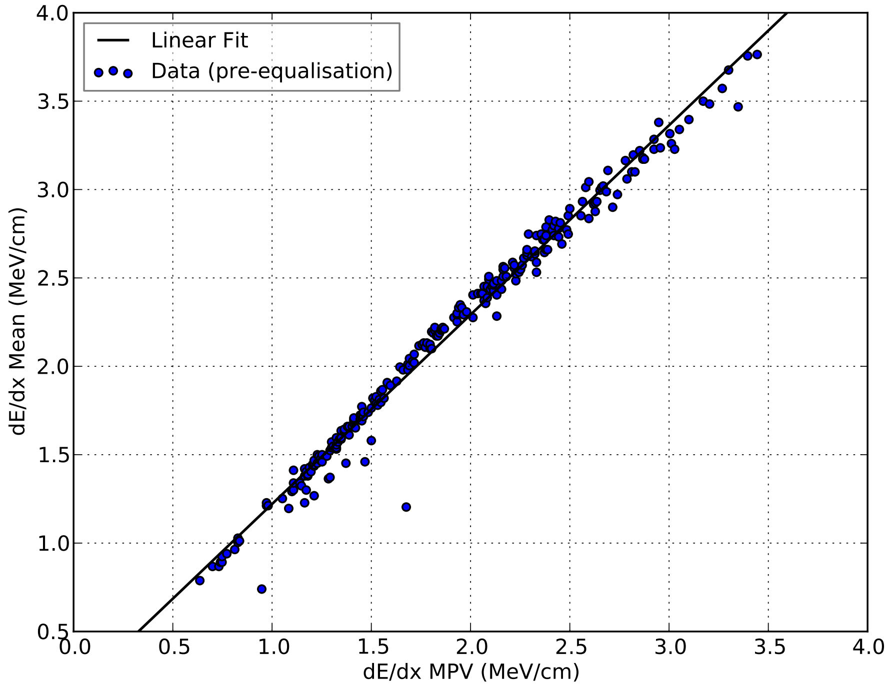

Figure 5.11: Correlation between dE/dx mean value and fitted MPV for each SiPM channel. The outliers correspond to cases where the fitting procedure failed.

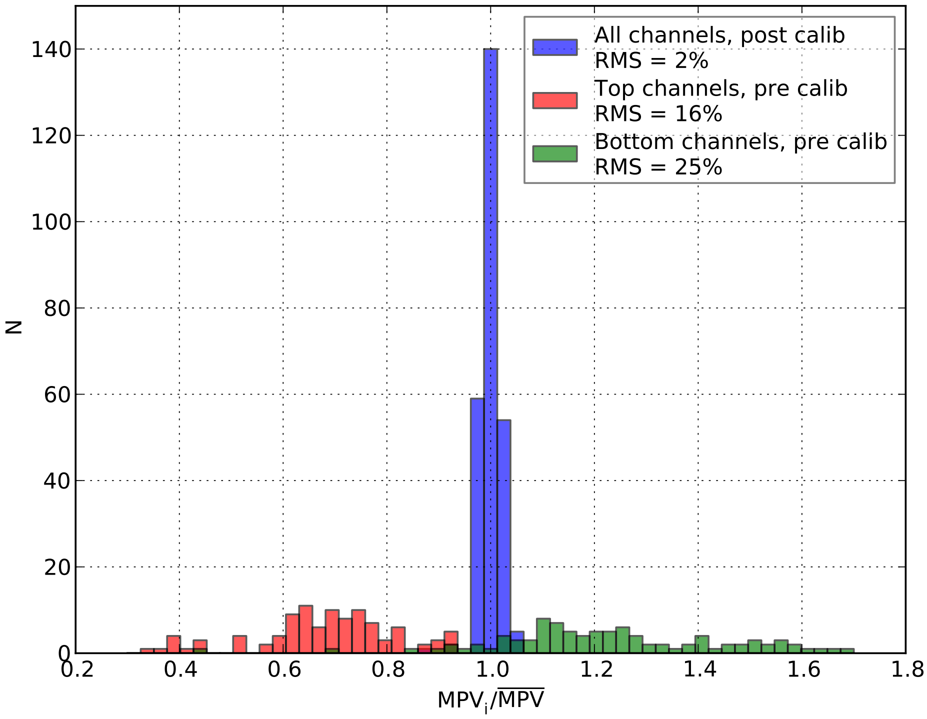

Figure 5.12: Pull distribution of fitted channel MPV, before and after (using separated data sets) the application channel equalisation.

trend can be observed for cubes numbered 0 to 140, and the opposite trend is observed for cubes numbered 170 to 300, corresponding to fibres where the the SiPMs are attached to the opposite end of the fibre.

Attenuation of the fibres has previously been measured and a correction model found from fundamental principles. Recall from chapter 4 that the fractional light loss F_{Loss} at position x along the fibre can be modelled as the sum of two exponentials: one for the contribution of light directly from the cube to the SiPM, and another from the reflection of the mirror at the other end of the fibre:

$$F_{Loss}(x) = \frac{1}{2}e^{-x/a} + \frac{1}{2}re^{-(2X-x)/a} \tag{5.2}$$

where X is the total length of the fibre, a is the attenuation length, and r is a measure of the reflectivity of the mirror (bound between 0 and 1, and assumed to be channel independent). The current measurements from lab setups are:

$$a = 1.3 \pm 0.1\text{m}^{-1} \tag{5.3}$$

$$r = 0.85 \pm 0.05 \tag{5.4}$$

In the extreme case when comparing one end of the fibre to the other, this results in around a 20% loss of light. The model is compared with the data from the prototype in figure 5.14, which shows the MPV and dE/dx as a function of position along the fibre (integrated over the equalised channels). Excluding the outer cubes, the data and model are reasonably consistent. The light yield of the outer cubes appear reduced. This has been understood as differences due to the reflectivity of the cube wrapping - cubes in the bulk of the detector are surrounded by other cubes, whose wrappings act as an additional layer of wrapping for the central cube. For cubes at the end of the fibres, at least one face is exposed to a different material (black HPDE), which is less reflective, thus lowering the light yield of the outer cubes.

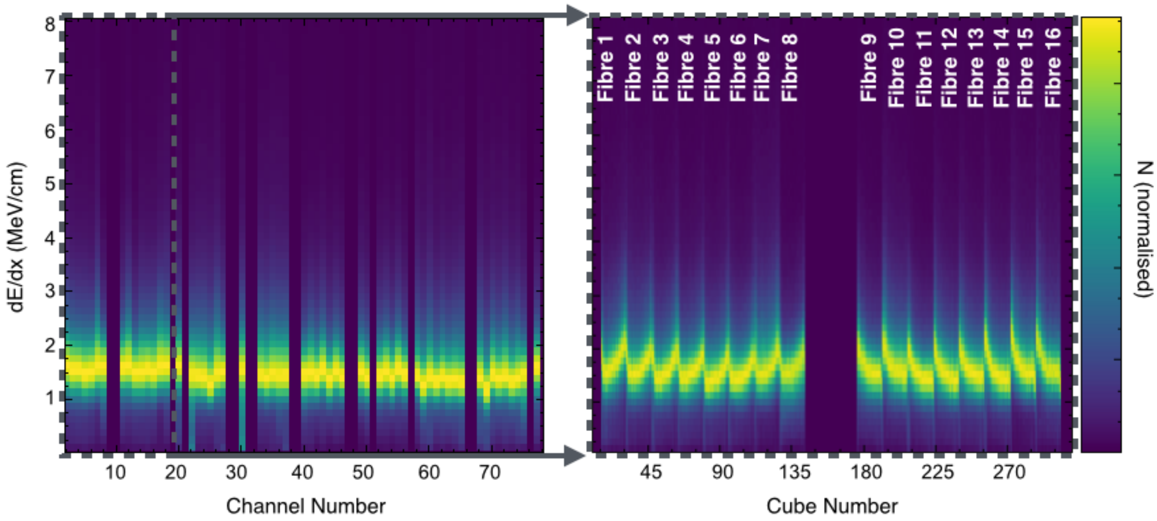

Figure 5.13: Diagram showing a zoom of the per channel dE/dx, revealing attenuation effects at the cube level. For this figure, cube number is 16 $\times$ channel number + position along the fibre. The empty slices correspond to masked channels.

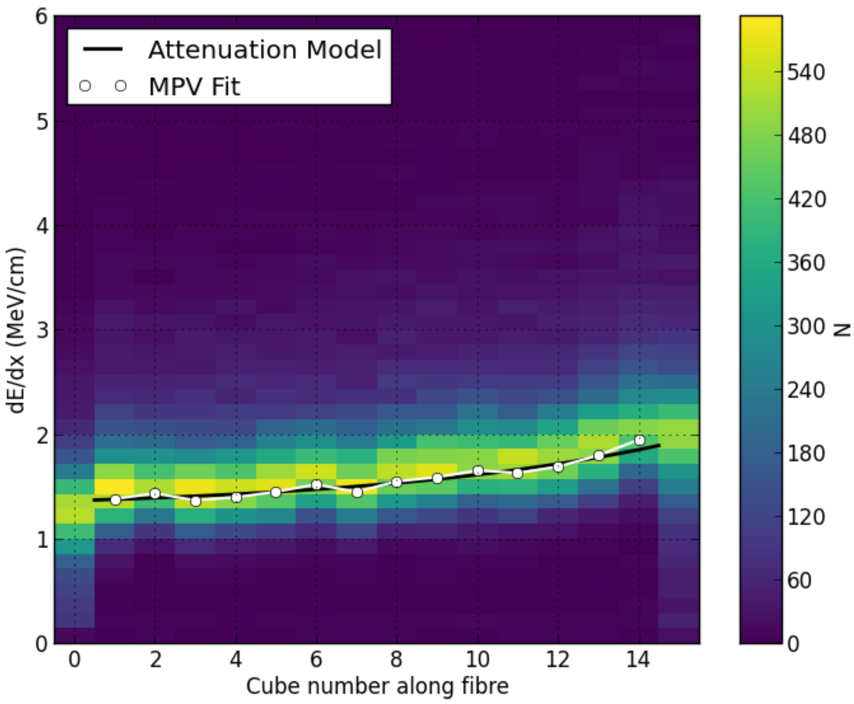

Figure 5.14: dE/dx as a function of cube position along fibre (integrated over the equalised channels). The data points show the MPV of each slice, and the black line shows the attenuation model prediction.

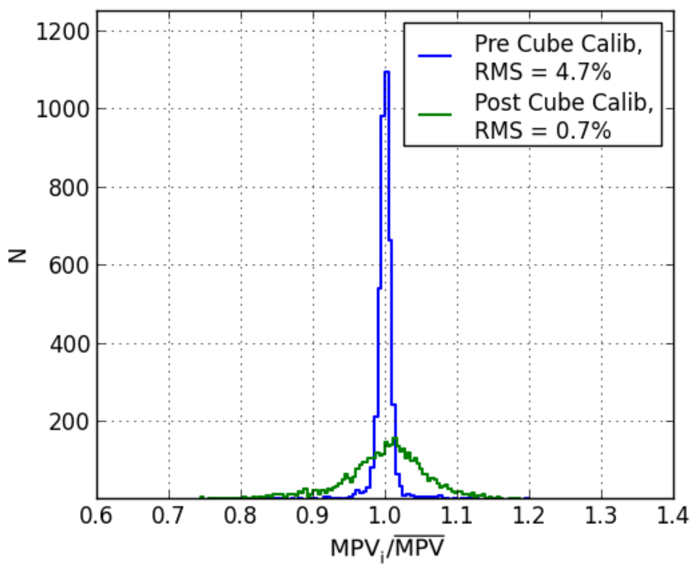

Figure 5.15: Pull distribution for fitted MPV for cubes, before and after offline cube equalisation.

5.4.3 Cube Equalisation (Experimental)

The above equalisation procedure used for the detector channels can be repeated to equalise any remaining cube-to-cube variations in light yield response. These can stem from variations in the coupling between the cubes and fibres themselves, which has been previously observed in a lab setup. The pull distributions when performing equalisation at the cube level are shown in figure 5.15. The spread before equalisation (with channel calibration and attenuation correction enabled) is 5%, suggesting there are significant light yield variations not caused by channel or attenuation effects. Post equalisation, this can be reduced to 1%. There are still some outliers, and these cases are under further investigation.

Note: this cube equalisation (post attenuation correction) is not yet implemented in default data reconstruction, and so not used in physics analyses such as the IBD analysis presented in chapter 6.

5.4.4 Energy Scale

Using dE/dx also provides a way to extract the absolute energy scale (i.e. light yield) of the detector. The dE/dx distribution predicted from simulation, not including any read-out effects, is shown in figure 5.16. The distribution can be described by a Landau distribution,

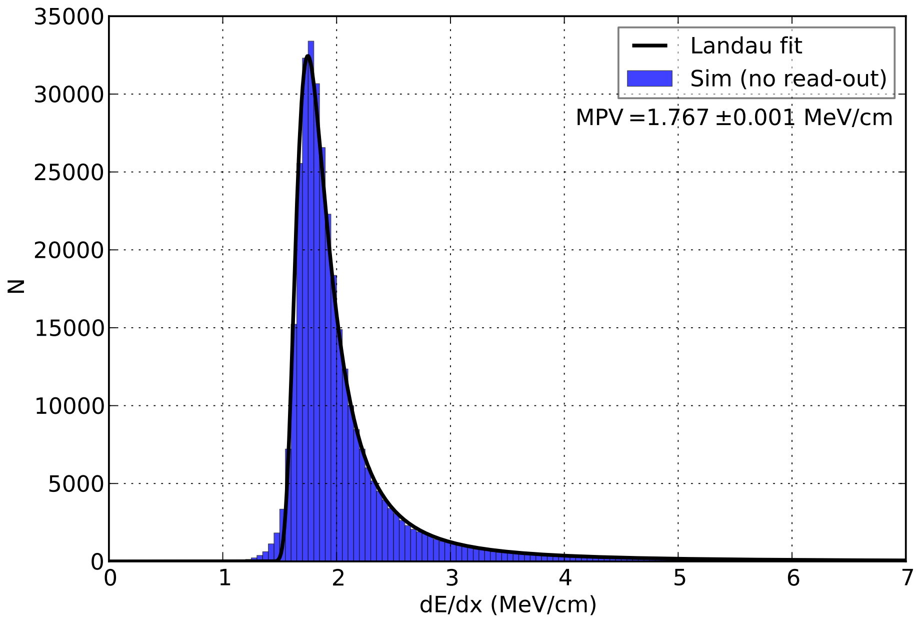

Figure 5.16: dE/dx for cosmic muon simulation (readout effects not simulated).

as expected, with an MPV of 1.767 ± 0.001 MeV/cm. The equalised detector data, in units of pixel avalanches, is linearly scaled to match the MPV value found here. The number of pixel avalanches PA per MeV (per fibre) is found to be:

$$1 \text{ MeV} = (12 \pm 1) \text{ PA} \tag{5.5}$$

The dominant error arises from the uncertainty of the gain measurement used to convert the digitised samples into pixel avalanches [45]. This measurement is consistent with other lab based measurements that are more precise: (11.2 ± 0.5) PA/MeV/fibre [63].

5.5 Simulation Validation

The equalisation procedure outlined above can been validated using simulation. This has been tested using cosmic simulations using Geant4 [64] with the Guan cosmic muon generator [69] (see chapter 7). Muons are propagated through the prototype module, including a full simulation of the geometry of the reactor building. The readout simulation is also enabled.

Figure 5.17 shows the distribution of MPV values (normalised to the expectation from

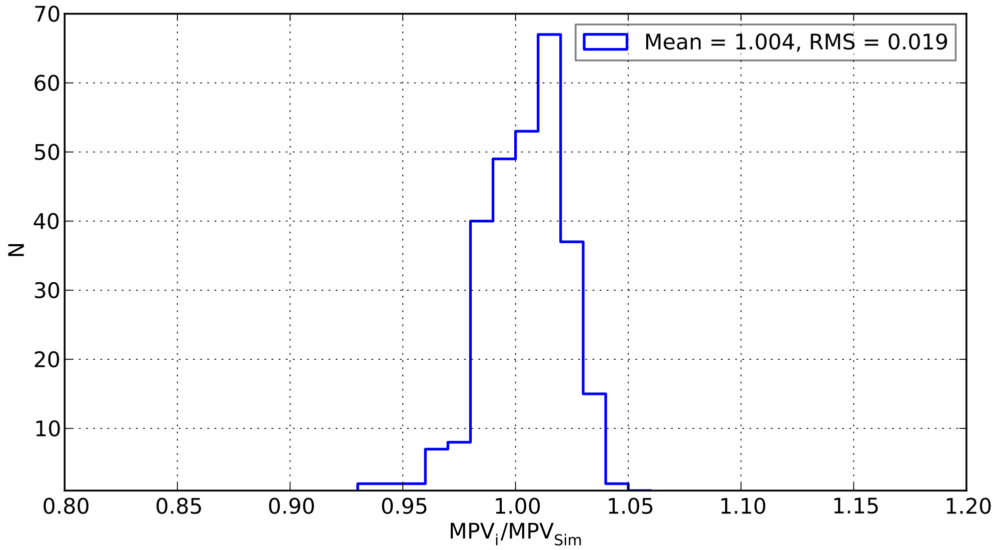

Figure 5.17: Spread of channel MPV for simulated cosmic muons (normalised).

section 5.4.4) over all channels using simulations of cosmic radiation. Since the distribution is centred around one, this suggests that the muon MPV is indeed a standard candle. The RMS of the distribution is 2%, suggesting that the assumption of a fixed MPV across all cubes is valid to the few percent level.

5.6 Conclusions

An energy calibration technique using cosmic muons has been developed for the SoLid prototype detector. Utilising the segmentation of the detector, a subset of muons that can be tracked to high accuracy as they pass through the detector has been identified. By imposing additional geometrical selection criteria, the energy deposits from these muons can be localised to individual cubes, allowing the dE/dx distribution to be found for each cube. Using this distribution as a standard candle, the energy response of the detector has been equalised to the cube level. This includes corrections for channel-to-channel variations, where the spread in channel response is reduced from 43% to 3%, and an attenuation correction due to light lost in the optical fibres. The remaining variations are at the scale of $\sim 5\%$ across cubes, which is an acceptable spread given the energy resolution of the detector. The validity of this approach has been verified in detailed simulation studies. Other methods to improve this further (such as per cube corrections) are being investigated for phase 1.

From comparisons with simulation, the absolute energy scale (i.e. light yield) has been measured to 8% accuracy, and is consistent with lab based measurements. The dominant uncertainty in this measurement arises from the measurements for the SiPM gains, which are challenging due to noise present in the waveforms. The spread in dE/dx MPVs using simulated data was also around 2%, placing a limit on the precision of the method.

The next chapter describes how calibrated events are used to search for neutrino candidates in the prototype data.

Chapter 6

Inverse Beta Decay Search at the Prototype SoLid Module One

6.1 Introduction

As well as a long reactor off run, prototype SoLid Module One also recorded a short three day reactor on run. The main purpose of this analysis is to have a thorough understanding of the detector, the signal and background shapes, and to search for the neutrino signal that may be present if these backgrounds are small enough. This is made more challenging given the reduced neutron detection efficiency, and the lack of passive shielding. The search is performed by defining a selection to isolate a region of phase space with increased signal-to-noise ratio, to give the highest chances of observing the neutrino signal. The segmented design allows for new and previously unexplored dimensions to this phase space, and the resulting reduction of the background is quantified here using data for the first time. Several techniques exist to perform such a selection, many of which benefit from the signal free (i.e. reactor off) background data sets that are available when performing reactor based experiments.

The chapter is structured as follows. Section 6.2 qualitatively describes the IBD signal and backgrounds. The reconstruction of IBD candidates is outlined in section 6.3, and various techniques are then used to perform a low background selection in sections 6.4 and 6.5. The chapter ends by applying these selection criteria to the reactor on dataset to search for the neutrino signal. This includes a prediction of the signal rate and a discussion of the statistical significance of the results.

6.2 Signal and Backgrounds

6.2.1 Signal Simulation

The neutrino signal is studied using many software packages. The simulation of the reactor core, and the propagation of neutrinos outward into the detector, uses the packages MCNPX and CINDER90 [57]. MCNPX is used to describe the geometry of the BR2 building, and CINDER90 calculates the fuel evolution of the reactor. To predict the neutrino energy distribution, the output of each type of fission reaction is weighted by predictions of the neutrino energy spectrum, provided by Huber et. al. (such as those found in [36]). Using this procedure, an estimate of the rate of neutrino interactions with the prototype module for the February 2015 reactor cycle has been calculated, and is shown in figure 6.1. The total number of predicted interactions for the 50.9 hours run is (prior to all other selections, and before taking into account the trigger efficiencies):

$$N_{IBD\ Interactions} = 630 \pm 48 \tag{6.1}$$

The error is dominated by errors associated with the power production of the reactor, such as calibrations of external detectors (this is incorporated as a systematic error below). The generation of IBD interactions in the detector is based on the formulae provided by Vogel [65, 71]. For each IBD interaction, the positron scintillation and annihalation, as well as the neutron thermalisation, are simulated using Geant4 [64]. The effects of the readout system of the prototype are also incorporated using an additional process, the 'read-out simulation', which includes trigger emulation and modelling the response of the SiPMs. This part of the simulation has to be tuned to data, and the errors associated with this tuning are also incorporated as systematic errors. The simulated data is then passed through the reconstruction chain (using Saffron1), just like real data, as outlined in the previous chapter.

The masking of detector channels is also simulated. The fraction of active cubes per detector plane of the large scale prototype is shown in figure 6.2. Channels were masked during online data taking due to especially large levels of electronics noise, resulting in a high number of additional triggers. Since an IBD candidate is formed of two triggers, this coincidence requirement will reduce the detection efficiency further (compared to studying single events). This has been quantified using simulation, and it is found that the fraction of IBD interactions

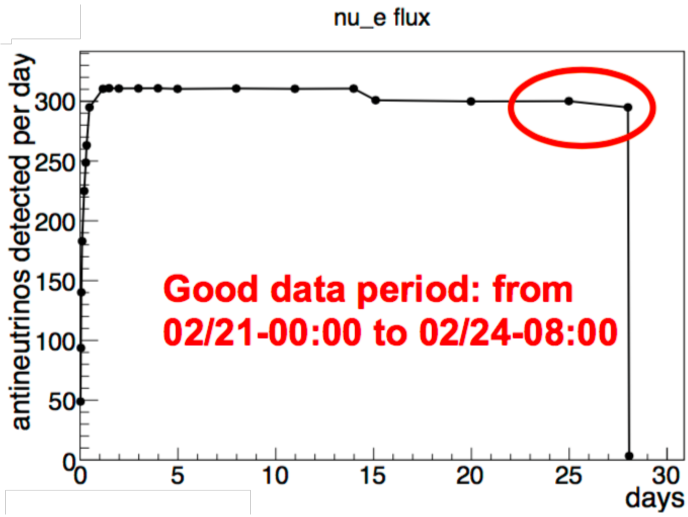

Figure 6.1: Predicted neutrino interaction rate at the large scale SoLid prototype module for the February 2015 reactor cycle. This is prior to any detector and analysis effects (such as trigger and analysis cuts).

not detected due to masked channels is 11%. The effect of masked channels, combined with channel dependent trigger thresholds (also included in the simulation), results in a significant variation in the neutron detection efficiency as a function of position in the detector. This is demonstrated in figure 6.3 using simulations of IBD interactions. As a result of such variations, some shapes of the IBD position distributions are distorted.

6.2.2 Background

The background of the IBD signal can be studied using data driven techniques. It can be separated into two categories: accidental and correlated.

Accidental Background

SoLid events that are not correlated with other events are often referred to as singles, and types of this event include reactor gamma rays and muons. Given the large rate of these single events, there is a high probability that a ZnS single and EM single are in coincidence at the neutron thermalisation time-scale, giving rise to IBD candidates that are 'mismatches' or 'accidentals'. This gives rise to the accidental background: B_{Acc}. The sources of this background can be sub-categorised further:

- $B_{Acc,\, Environment}$: collective term used for all sources of accidental background when the reactor is off. This can include muons crossing a small part of the detector, which can

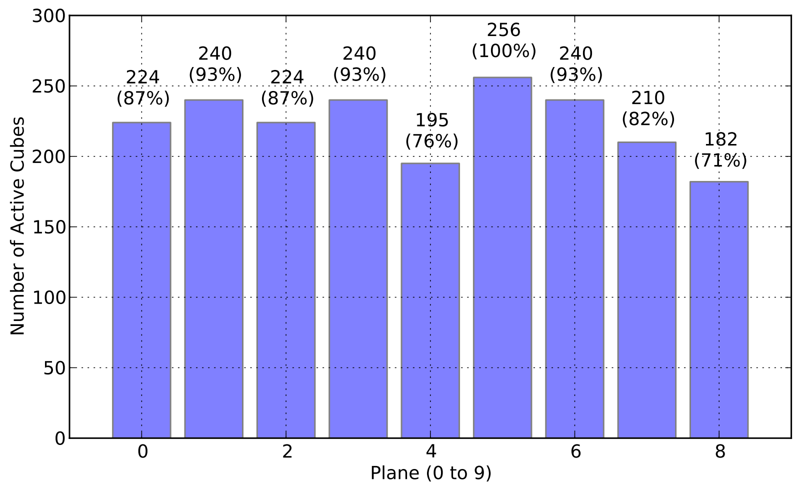

Figure 6.2: Number of active cubes per plane, as a result of masking SiPMs duing commissioning. If a SiPM is masked, the corresponding row of cubes is inactive.

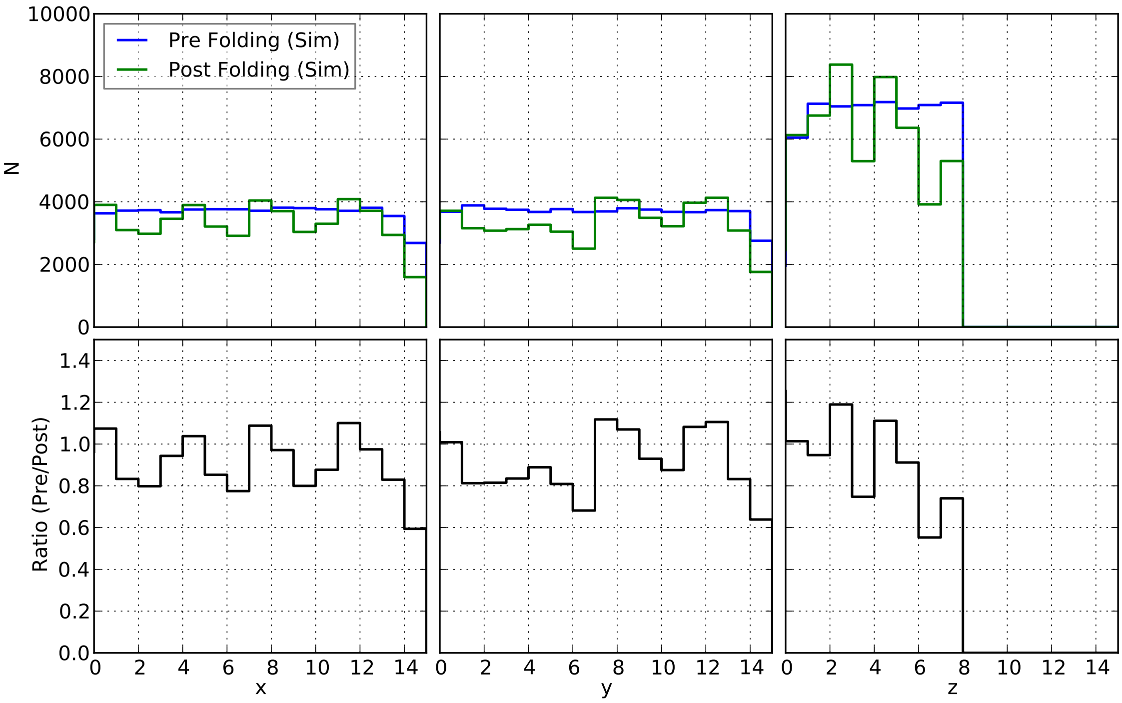

Figure 6.3: Neutron position distribution for the IBD signal simulation, before and after applying masked channels and varying trigger threshold across channels.

be difficult to identify from positions, as well as radioactive decays from materials other than the reactor.

- $B_{Acc,\,Reactor}$: the additional contribution to the accidental background caused by the increased rate of single events when the reactor is operational, such as reactor gamma rays. This is defined as the difference between the accidental background for reactor on and off.[1]

Topologically, the positron candidate (AKA the prompt, P) and neutron candidate (AKA the delayed, D) events forming the accidental backgrounds are typically well separated in space. This is expected, since these uncorrelated events occur randomly in the detector volume, and so their rate increases with volume. As such, the granularity of SoLid is well placed to combat this type of background. Further, the energy distribution of the accidental background positron candidates is shifted lower than that of the signal, and the the time difference between the prompt and delayed events is also typically much larger. These variables can be used to discriminate the accidental backgrounds from the signal (distributions shown later).

The shapes of the accidental background are studied using shifted time windows. Specifically, the timestamp of neutrons is shifted in software forwards in time by 1 ms, forcibly removing any correlations between EM and ZnS events at the neutron thermalisation time-scale. This is illustrated in figure 6.4. This provides a dataset guaranteed to be formed of only accidental candidates. Due to the increased accidental background when the reactor is powered, this is done for both the reactor on and off datasets separately:

$$B_{Acc,\,On} = B_{Acc,\,Environment} + B_{Acc,\,Reactor} \tag{6.2}$$

$$B_{Acc,\,Off} = B_{Acc,\,Environment} \tag{6.3}$$

Correlated Background

The correlated background B_{Cor} is formed of ZnS and EM events that are produced via the same mechanism, which itself is not a neutrino interaction. The dominant contribution

[1] Since accidental backgrounds increase quadratically with the single rate, this term therefore does not describe the accidental background caused by the reactor only (i.e were $B_{Acc,\,Environment}$ zero). This definition is used in order to study the transition between reactor on and off in more detail.

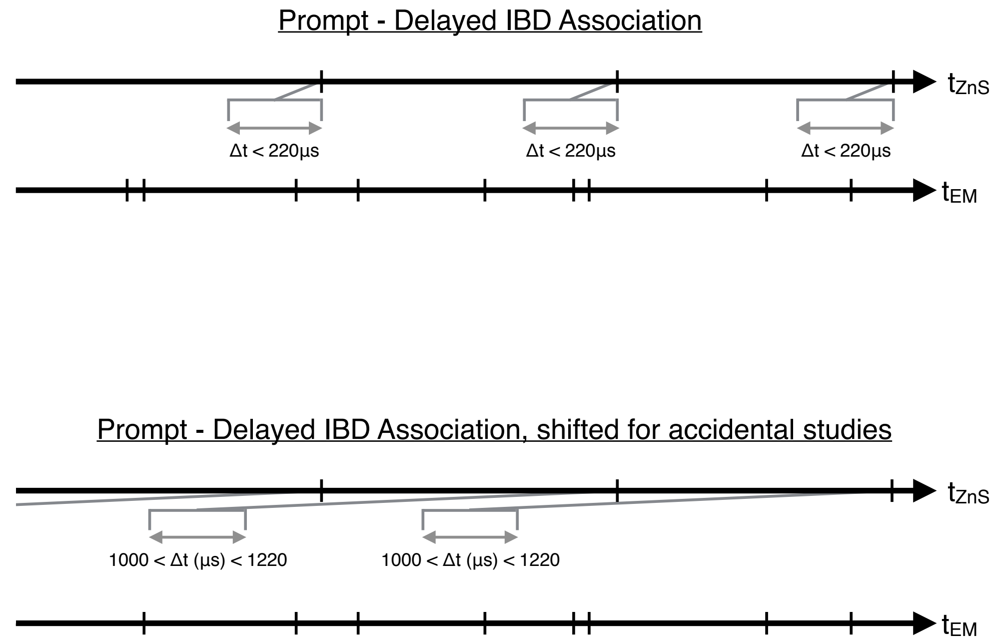

Figure 6.4: Illustration of using shifted time windows to study accidental backgrounds. The marks on the time-lines represent events. Top: example time-lines for ZnS and EM events, where the gray box shows the association time window for forming IBD candidates. Bottom: the same but using a shifted time window.

to B_{Cor} is thought to be primarily from fast neutrons - highly energetic neutrons, which during their thermalisation, induce proton recoils by elastic scattering [41]. The energy of these protons is mostly <100keV (the distribution is shown in chapter 7), but can extend to much higher energies $\mathcal{O}(10)$MeV. Some protons can therefore mimic a positron signal. Since the time difference between the prompt and delayed events is determined by the neutron thermalisation time, the time difference is similar to that of neutrino signal. Topologically, the fast neutron background also behaves similarly to IBDs, although the proton recoils can have higher multiplicity (see below). As a result of these similarities, the correlated background is more difficult to reduce.

The primary source of fast neutrons at SoLid are from cosmic origins. Interactions between cosmic rays (primarily protons) and nuclei in the atmosphere can produce secondary showers. The products of these interactions include fast neutrons and charged mesons, such as pions [42]. The primary decay of pions is to muons, which in turn, can also induce fast neutrons via muon spallation whilst en route to the detector (e.g spallation with the reactor hall) [43]. Muons themselves can also lead to correlated background events, if a muon spallation interaction occurs within the detector itself.

The correlated background is studied using the reactor off dataset, with the reactor off accidental background subtracted. Given the reactor is not a source of fast neutrons, the correlated background rate is expected to be independent of reactor power.

Model

Using the notation introduced above, the total expected background for the reactor on period B_{On} can be written as:

$$B_{On} = B_{Acc,\,Reactor} + B_{Off} = B_{Acc,\,Reactor} + B_{Acc,\,Environmental} + B_{Cor} \tag{6.4}$$

Since the reactor off dataset is relatively large, the rate and shapes of these background contributions can be found to high precision. The background rates are scaled to the duration of the reactor on period (see section 6.6).

6.3 Inverse Beta Decay Reconstruction

IBD interactions are formed of an association between an EM event and a ZnS event. Forming this association not only allows for a significant reduction in background compared to studying just individual events, but also gives access to further discriminating variables, such as the spatial and temporal separation between the prompt and delayed events.

6.3.1 Prompt-Delayed Association

The association is performed by searching for time correlations between events at the neutron thermalisation time-scale $\mathcal{O}(100)\mu$s. This is done by placing selection criteria on the time difference between ZnS (t_D) and EM (t_P) events: $\Delta t = t_D - t_P$. The algorithm is implemented as a double loop:

- Loop over all ZnS events (ordered in time):
  - Loop overall EM events (ordered in time):
    - If the time difference between the delayed and prompt event Δt is less than some time-window, form an IBD candidate.

As seen, no rules are used to prevent one event from forming multiple IBD candidates, even though this would be non-physical. The resulting 'extra' candidates would be due to mismatches between prompt and delayed events, and therefore contribute to the accidental background. Whereas extra conditions could be used to skip events that already form an IBD candidate, or perhaps break the EM loop once a candidate is formed, such rules could mean that actual IBD interactions may not be reconstructed at all. It is found that post all selections used to focus on the IBD signal, the fraction of IBD candidates that arise from multiple prompt events per delayed is less than 1% of the total background, and therefore negligible.

Δt Selection

The Δt distribution is shown in figure 6.5 for the reactor off dataset, prior to any other selections. This plot shows many features:

- A flat component is visible around N = 1400, for both positive and negative Δt, corresponding to the accidental background.
- A large correlated component at positive Δt, which corresponds to correlations between events (i.e. the correlated background).
- A smaller correlated component at negative Δt. This effect has been observed by multiple experiments [50], and corresponds to cosmic neutron showers. These showers can involve multiple fast neutrons. This contribution is interpreted as cases where neutrons are incorrectly correlated with prompts produced by other neutrons from the same shower. In principle, these could be removed by applying a multiplicity cut on number of neutrons in time coincidence; however, due to the low neutron efficiency, it is found this has negligible effect.

The value used for the prompt-delayed association time-window is $\Delta t < 220\mu$s, and is set to contain 91% of correlated candidates (found from simulation - errors discussed later). This fraction is consistent with the values set for all other IBD cut-based selections that are described in section 6.4.

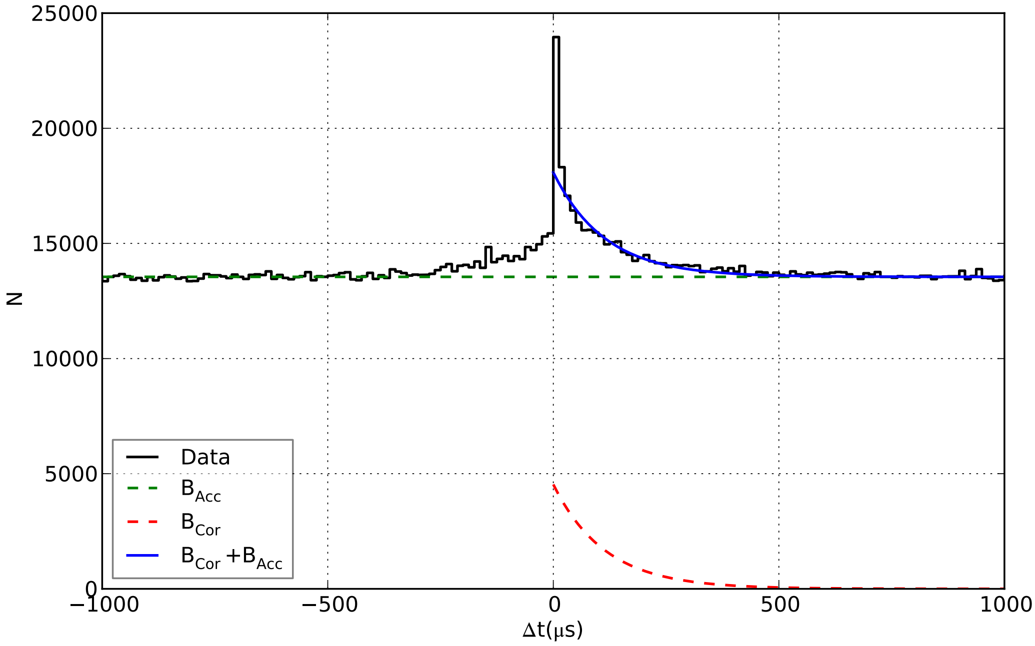

Figure 6.5: Distribution of the time difference between prompt and delayed events for IBD candidates, prior to any IBD selections. An exponential fit to the correlated component is plotted, although as this is not used in any analysis presented here, this is only for visualisation.

Since Δt is used at this earlier reconstruction stage to initially form IBD candidates, as well as for IBD selections described later, it has special status amongst the other IBD variables. Further, as some IBD variables can only be extracted after this association, some results (such as selection efficiencies) are made relative to this reconstruction stage. This special status is however just a convention - other conventions could be used, and this has no bearing on the final results.

Using this association criteria, the fraction of ZnS events forming IBD candidates is 61.1%, and the fraction of EM events forming IBD candidates is 1.3%. As can be seen, as well as giving access to new discriminating variables, using this co-incidence requirement already gives a dramatic reduction in backgrounds compared to considering just EM events in isolation. This is demonstrated by revisiting the EM event energy distributions, which has been re-plotted in figure 6.6. It can be seen the number of EM prompt candidates is reduced by around two orders of magnitude (comparing the green solid line with the green dashed

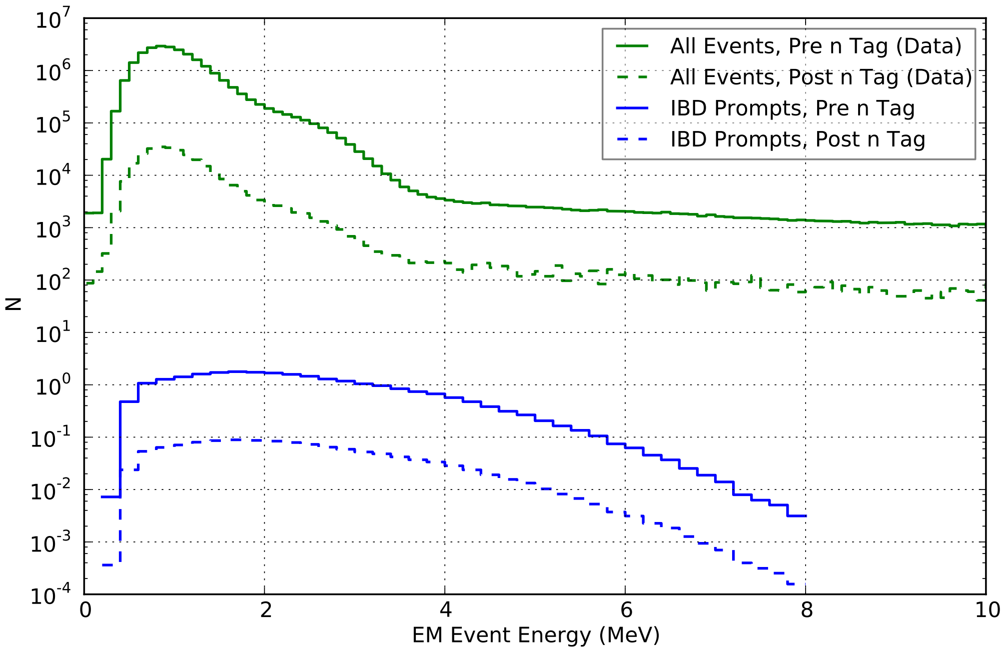

Figure 6.6: The energy distribution of EM events (i.e singles), before and after association with neutrons. The IBD results are form simulation, whereas the other curves are found using data.

line). However, this reduction in background comes at a significant cost of signal efficiency of a factor ~ 30, due to the low neutron detection efficiency of the prototype.

6.3.2 IBD Feature Extraction

The following features are extracted for each IBD candidate:

- Δt: time difference between the D and P events, $\Delta t = t_D - t_P$.
- Topological:
  - Δz: separation between D and P in the z direction (recall z points away from the reactor), $\Delta z = z_D - z_P$.
  - Δxy: radial separation between P and D in the xy plane, $\Delta xy = +\sqrt{(x_D - x_P)^2 + (y_D - y_P)^2}$. Δxy is used instead of using Δx and Δy separately due to the symmetry in the xy plane.

These are in addition to the features of each of the two events forming the IBD candidate, such as positron volume: V_{Prompt} (recall that volume is a measure of multiplicity). The distributions of all these features for backgrounds (data) and signal (simulation) are shown throughout sections 6.4 and 6.5.

6.4 Cut-Based Selections

This section investigates the effects of simple hand chosen cuts placed on the IBD features to reduce backgrounds. This includes studying the shapes of the signal and background distributions, with the ultimate aim of optimising the signal-to-noise ratio (for a given signal efficiency). This is a useful first step to reducing the background in the analysis, and more sophisticated techniques are considered in the next section.

The selection criteria themselves, for a given variable, are typically set to contain approximately 90% of signal events, to retain a sufficiently high fraction of the signal. This is set prior to placing any other cuts on the other variables. The variables themselves include: Δxy, Δz, V_{Prompt}, and E_{Prompt} (i.e prompt energy). The first three are referred to as the topological cuts.

6.4.1 Plot Style

For each cut placed on a variable, the distribution is presented for an extended range available (i.e. just beyond the cut), and in two ways: *prior* to all other cuts, and *post* all other cuts. The three categories of background plotted are: B_{Cor}, $B_{Acc,\,Reactor}$ and $B_{Acc,\,Environment}$ (not stacked). Their contributions are scaled to the expectation for the reactor on period. The signal expectation from simulation (including readout and reconstruction) is also plotted, and scaled upwards by 100$\times$ the expected rate for clarity. Further, the normalised cumulative distributions are also plotted, to show selection efficiencies. At the end of the section, the signal efficiency and background reduction of all cuts will be summarised.

6.4.2 E_{Prompt}

The prompt energy distribution of IBD candidates is shown in figure 6.7. Many effects can be seen. The dominant source of background, both before and after the application of the cut-based selections, is the accidental background from the reactor. The energy distribution of all background components is shifted to smaller values compared to the IBD signal. This includes the correlated background, which can also be seen to extend to high energies (well above 8MeV), prior to other applied selections. This is most likely the muon contribution, as post other selections, this contribution is significantly reduced (recall that the volume and energy variables are correlated for muons).

The default method for placing a selection is to include $\sim 90\%$ of signal events. However, for this variable, given the large amount of background at low energy and the steepness of its distribution, a harsher selection is placed at 1.5 MeV to increase the signal-to-background ratio. This has a selection efficiency of 73%, and reduces the remaining background post other cuts by a factor of 4.0. An upper selection is also placed on energy at 8 MeV, which is 100% efficient for signal. This upper cut is effective for reducing the correlated background prior to other cuts, although as seen, post other cuts this is redundant due to correlations between the variables for this background. In this region, the correlated background is the dominant remaining contribution.

6.4.3 Topological

Δxy

The spatial separation between the prompt and delayed events in the xy plane is shown in figure 6.8. The separation for signal is typically much smaller than that of the backgrounds, with $\sim 90\%$ of signal IBDs separated by 2 or fewer cubes (in the xy plane). All sources of background are distributed at higher values, with around 90% of accidentals above this region, both before and after the application of other cuts. The shapes of the two accidental backgrounds are very similar - this is due to symmetry in the xy plane (i.e along the reactor axis). The majority of correlated events are also at higher separations, although the fraction decreases after applying other selections. These results demonstrate one of the advantages of the segmentation of the detector, especially for reducing accidental backgrounds; a factor of ten in background reduction can be gained post all other selections using this variable alone.

The default selection is $\Delta xy \leq 2$ cubes, which has a signal selection efficiency of 90%, and reduces the total remaining background by a factor of 5.7 - the largest of all the variables.

Δz

As well as allowing for measures of spatial separation, the segmentation of the detector also allows the directionality of events to be measured. Specifically, the distribution of $\Delta z =$

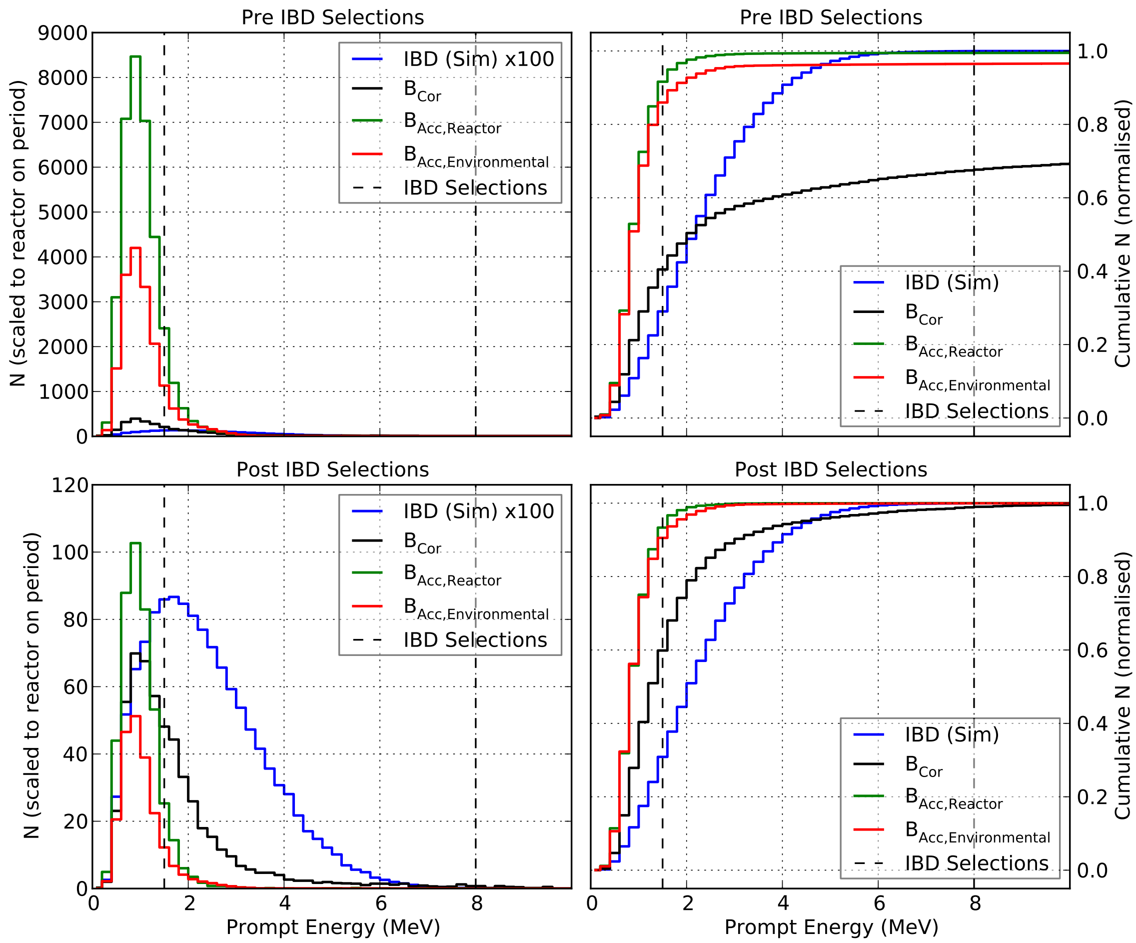

Figure 6.7: The energy distribution of IBD prompt events from simulation, and IBD backgrounds found using data, before (top) and after (bottom) cut-based selections. The plots on the right show the cumulative distributions (normalised). The range plotted is to focus on the signal region - the correlated background energy spectrum extends to larger energies.

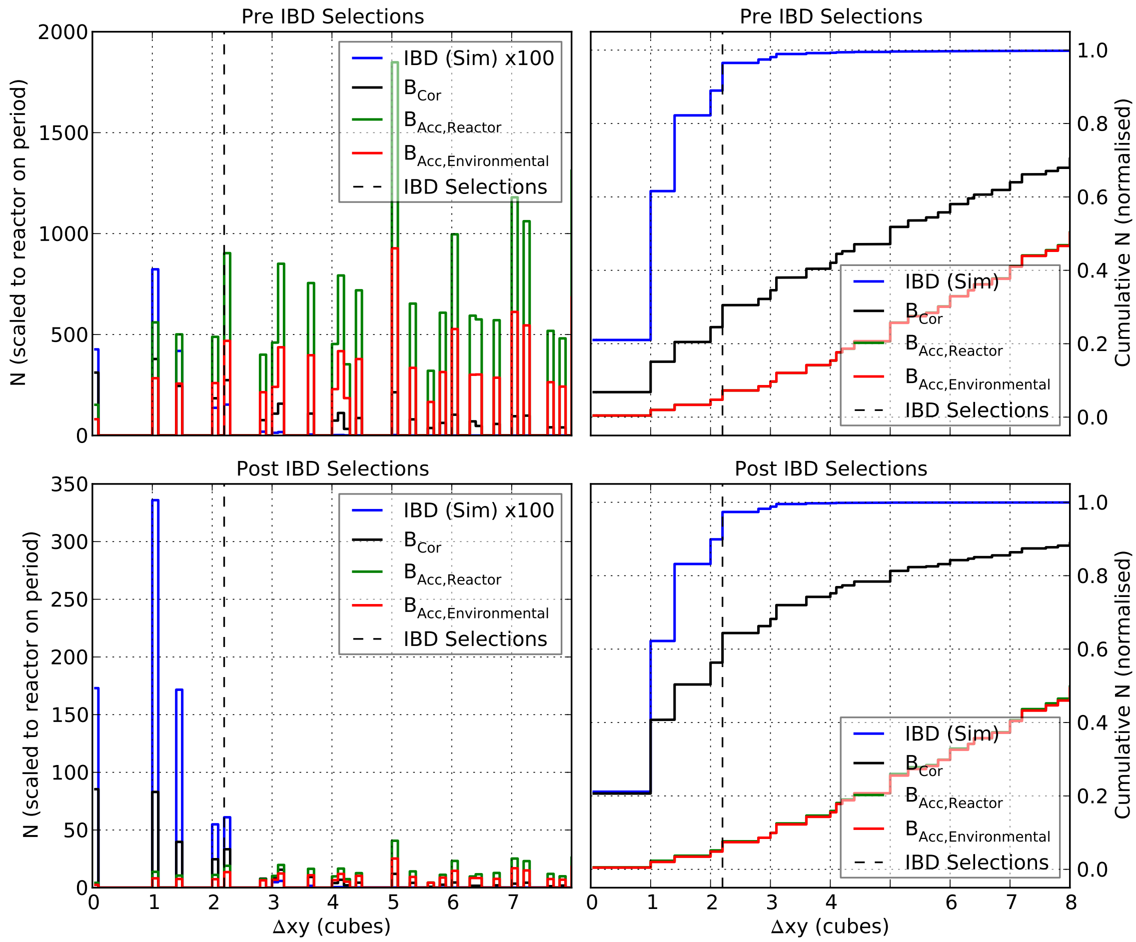

Figure 6.8: Radial separation between the neutron and positron events in the xy plane, before (top) and after (bottom) cut-based selections. Δxy is a discreet variable, giving rise to the spiky plots. Note that bins are left aligned (i.e not centred).

$z_D - z_P$ is expected to be shifted towards positive values for signal, as seen in figure 6.10. These correspond to cases where the positron is closer to the reactor than the neutron (on average), as in the right event display shown in figure 6.12. There are two effects that can contribute to this bias:

- **Momentum kinematics**: for signal IBDs, the neutron carries away the majority of the momentum, giving a small bias $\sim 2\,\mathrm{cm}$ in the direction away from the reactor - see [46].
- **Lithium screen position**: recall that the position of the ^{6}LiF:ZnS(Ag) is placed on the cube face that is closest to the reactor. The neutron capture position is therefore actually (z_D-0.5) cubes, causing a bias in the distribution of Δz to positive values.

The accidental backgrounds are more homogeneous. In addition to directionality, Δz also gives a complimentary measure of spatial separation to Δxy, which can also be used to remove $\sim 70\%$ of the remaining accidental events post other selections. The default selection is: $0 \leq \Delta z \leq 2$ cubes, which has a signal efficiency of 94%, and this reduces the total remaining background by a factor of 1.9.

Δr

For completeness, figure 6.9 shows the cumulative distribution of Δr (before other IBD selections). Here, $(\Delta r)^2 = (\Delta xy)^2 + (\Delta z)^2$, and is an overall measure of the radial separation between the positron and neutron. The same patterns from above can be observed - both background categories can be reduced, especially the accidental. Selection criteria are not placed on the variable by default, since the directionality information is lost - instead, Δxy and Δz are used separately (as presented above). Nevertheless, this parameter well demonstrates how spatial separation can be a powerful measure to reduce IBD backgrounds.

Prompt Volume

The volume of prompt events is used to specifically discriminate signal from cosmic backgrounds (such as muons and proton recoils). In $\sim 95\%$ of cases, the positron deposits energy either in one cube or two cubes that share a face (i.e. $V_{Prompt} = 2$) - see figure 6.11. Both correlated and accidental backgrounds have around 20% of the remaining background

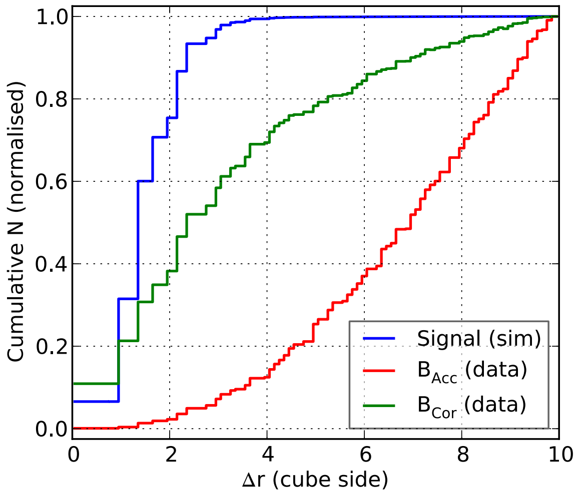

Figure 6.9: Cumulative distribution of Δr (the radial separation between the positron and neutron candidates), prior to other IBD selections. The accidental background is for reactor on data.

outside of this selection, and this fraction is much larger before the application of other selections.

The default selection in this variable is $V_{Prompt} \leq 2$ cubes. This has a signal efficiency of 94%, and removes 22% of the remaining background.

6.4.4 Event Displays

It is interesting to visualise candidates that have been classified using the above selection criteria. Two example candidates that fall within the signal region are shown in figure 6.12. The prompt and delayed events are clearly correlated in space, and the energy of the positron candidate is in the expected region (contained within a single cube). Example candidates that were rejected by the cut-based criteria are shown in figure 6.13. In both cases, both the volume and upper energy cuts reject these IBD candidates, suggesting the candidates are infact correlated background.

6.4.5 Discussion

Figure 6.14 summarises the efficiency of each selection, for signal and backgrounds, both before and after the application of all other selections. Each selection can be seen to be effective at background reduction in the absence of other cuts being applied. Post all other cuts applied, most variables still show significant background reductions, with the exception

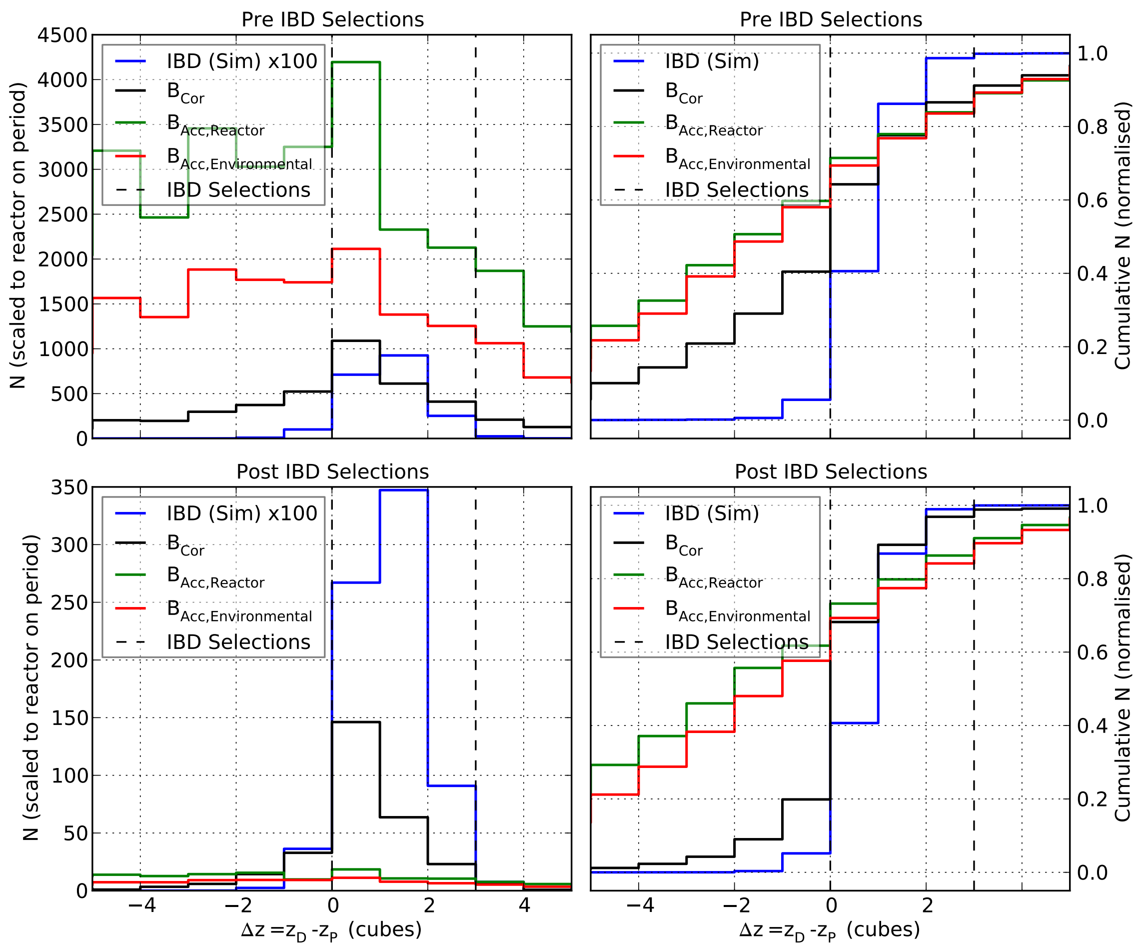

Figure 6.10: Spatial separation between the neutron and positron events in the z direction, before (top) and after (bottom) cut-based selections. For signal candidates, the positron is more likely to be closer to the reactor than the neutron.

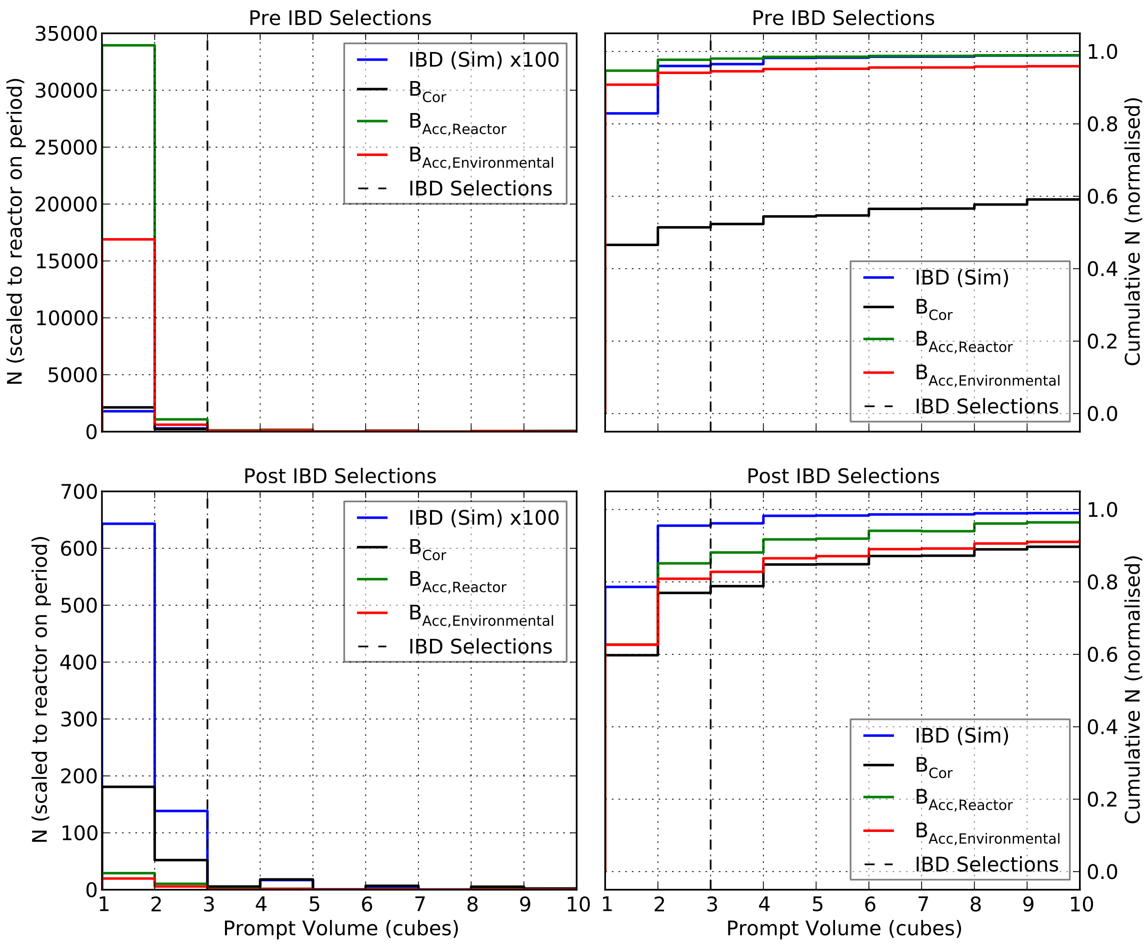

Figure 6.11: Distribution of prompt event volume, before (top) and after (bottom) cut-based selections.

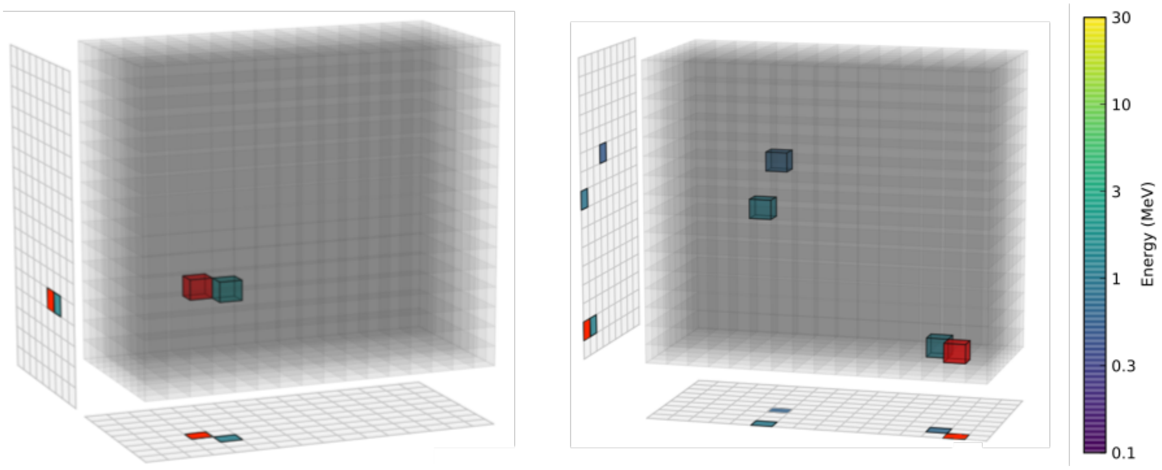

Figure 6.12: IBD candidate events from data that pass the cut-based selection criteria. The neutron events are marked in **red**, and otherwise an event is tagged as EM, and uses the colour scale to mark the energy of the event. For the event on the right, the two additional EM hits are not in fast time coincidence with any events, and happen to fall within the time window plotted ($\pm500\mu$s around the neutron) - given the high rate of gamma rays from the reactor, these coincidences are not unexpected. In both cases, the reactor is found behind the module (into the page).

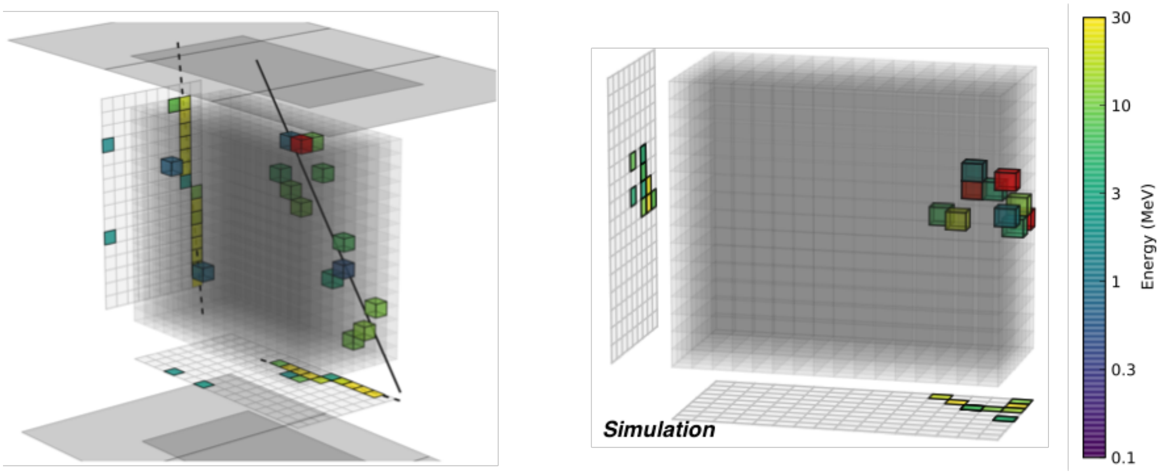

Figure 6.13: IBD candidates that do not pass the cut-based criteria. Left: example from data involving a muon track. Right: example from cosmic neutron simulation. The cube finding algorithm used to highlight specific cubes has not been optimised for these high multiplicity events, meaning some cubes have not been drawn, although this has no bearing on the final analysis. The colours use the same convention as used in figure 6.12.

of the upper energy cut (expected given these events are most likely muons, for which energy and volume are correlated). The correlated background is tackled most effectively by the volume selection, and this holds true post all other selections, suggesting a low energy high multiplicity component to the background (such as proton recoils or clipping muons). The accidental backgrounds are best tackled with the topological selections, specifically Δxy and Δz, and can be reduced by factors up to ~ 20 post other selections. Since Δxy is also effective at tackling correlated background, giving a reduction of factor ~ 2, this particular selection is the most effective of them all, reducing the total background by a factor of ~ 16 before other selections.

The overall reduction in background, both relative and absolute (scaled to the reactor on expectation), for each selection is shown in figure 6.15. This particular order of applying the cuts is used to focus on the effects of the topological variables. For an overall analysis selection efficiency of 57%, the total background is reduced by a factor of ~ 200. This value of efficiency is close to the target informed by sensitivity studies (see also chapter 7). The accidental backgrounds are significantly reduced by a factor ~ 1000. The correlated background is the remaining dominant background, having been reduced by a factor ~ 20. The signal-to-background expectation is $\sim 1:30$. On the benefits of segmentation, for the accidental backgrounds, segmentation provides an additional background reduction of over two orders of magnitude. The reduction is not so extreme for the correlated background, although a reduction of an order of magnitude is achievable. Finally, compared to the case of using single events (i.e. prior to the application of Δt), this demonstrates the signal-to-background ratio can be increased by a factor of ~ 350.

6.4.6 Multidimensional Distributions

The distribution of IBD features in four dimensions, for the most effective background reduction variables (Δt, E_{Prompt}, Δxy, and Δz), is shown in figures 6.16 and 6.17 for signal and background respectively[2]. Each individual scatter plot shows the 2D distribution of Δt vs. E_{Prompt}, and each plot corresponds to a bin of Δxy and Δz, which are both discrete variables (thus avoiding any projections over these variables). The region defined by the selections above is highlighted in blue. The 4D range has been chosen to focus just beyond

[2]Four dimensions are chosen due to the complications of visualising data in 5 dimensions, although that is possible, see [51].

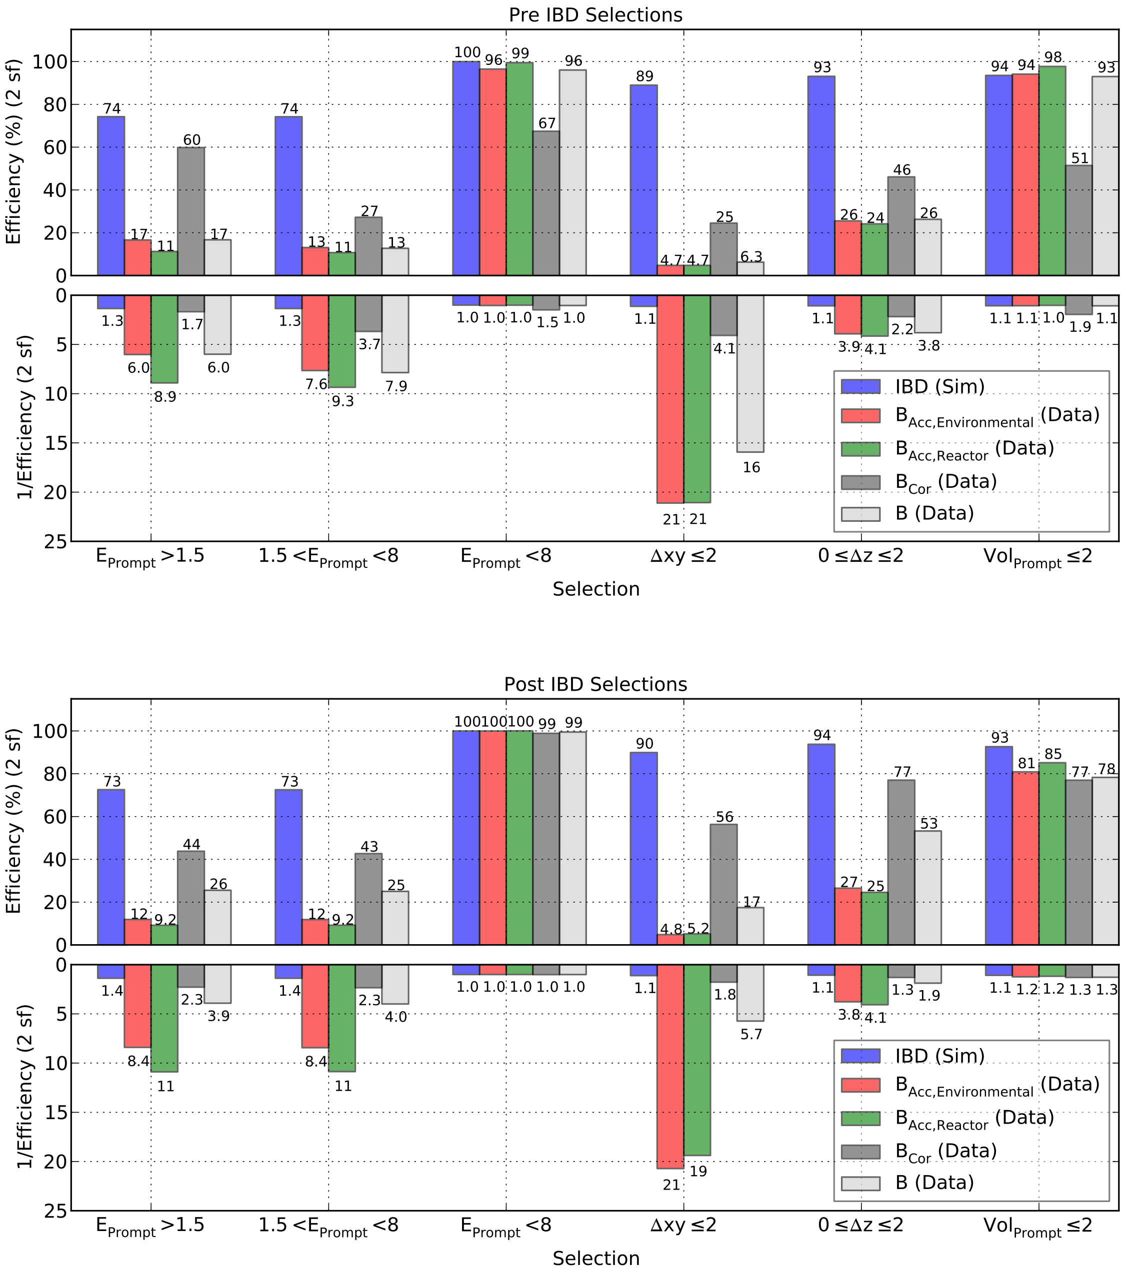

Figure 6.14: Selection efficiencies, before (top) and after (bottom) the application of cuts based selections. The inverse of these efficiencies is also included to show the reduction factor for each population.

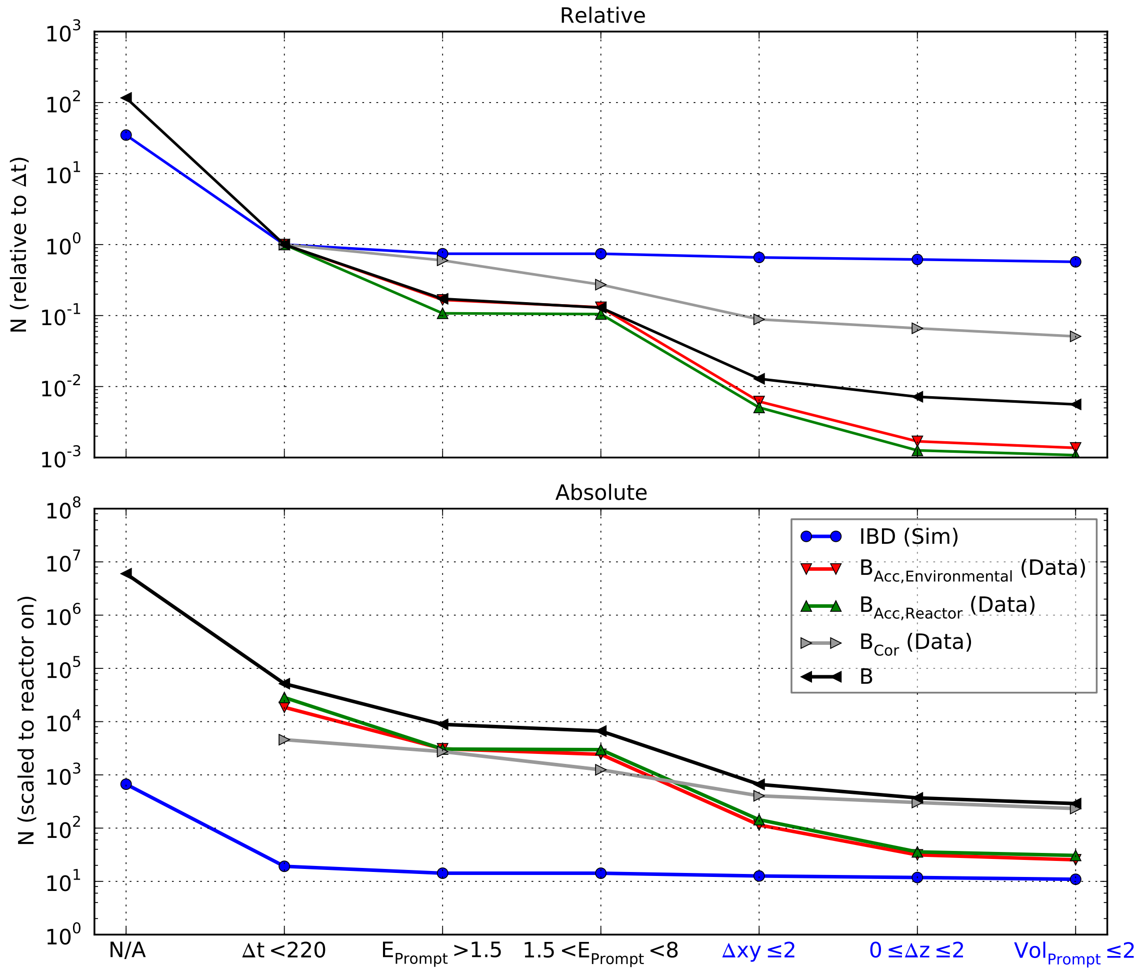

Figure 6.15: Signal and background rates for each selection cut applied sequentially. The order has been chosen such that selections on topological variables are applied last (highlighted in blue). The relative rates are normalised to the number of IBD candidates reconstructed (i.e. using a Δt cut only).

the signal region defined above - the background events extend well beyond the ranges shown (as seen with the cumulative distributions from above). This distribution is filled post IBD reconstruction - there are no additional cuts placed on any IBD variables.

The patterns observed are consistent with those described above: the signal is mostly contained at low values of Δt and in the energy range of 1MeV to 8MeV, for all bins of Δxy and Δz. Conversely, the background is spread over a larger range of Δt and at lower energies, as expected for the accidental background. The background candidates at larger energies most likely correspond to the correlated background, and many lie in a similar region to that of the signal. In the absence of other variables with which to place selections, this would be a form of irreducible background (i.e. background that cannot be removed by analysis without collecting more information about each IBD candidate, for example, by modifying the detector).

There are many signal IBD candidates that lie outside the signal defined rectangular region; and similarly, there are many background entries inside the signal region. The following section considers other techniques to re-define the signal region, with an aim to increase purity and efficiency of the selection than these rectangular cuts (shown in blue).

6.5 Multivariate Techniques

There are many other techniques that can be used to better separate signal and background. It can be analytically shown that in the case of classifying events into two categories, the best separation can be achieved by sampling the likelihood ratio formed of the two corresponding probability density functions P [59, 60]:

$$f(\mathbf{x_i}) = P_S(\mathbf{x_i})/P_B(\mathbf{x_i}) \tag{6.5}$$

$f(\mathbf{x_i})$ is also known as the Neynman-Pearson (NP) Classifier. $\mathbf{x_i}$ is a vector of measured variables of the ith event, P_S is the probability density function for the signal, and P_B is the probability density function of the background model. For signal and background classification at SoLid, each entry of x corresponds to an IBD candidate. In order to evaluate the NP-Classifier, P_S and P_B are required to be known for both signal and background prior

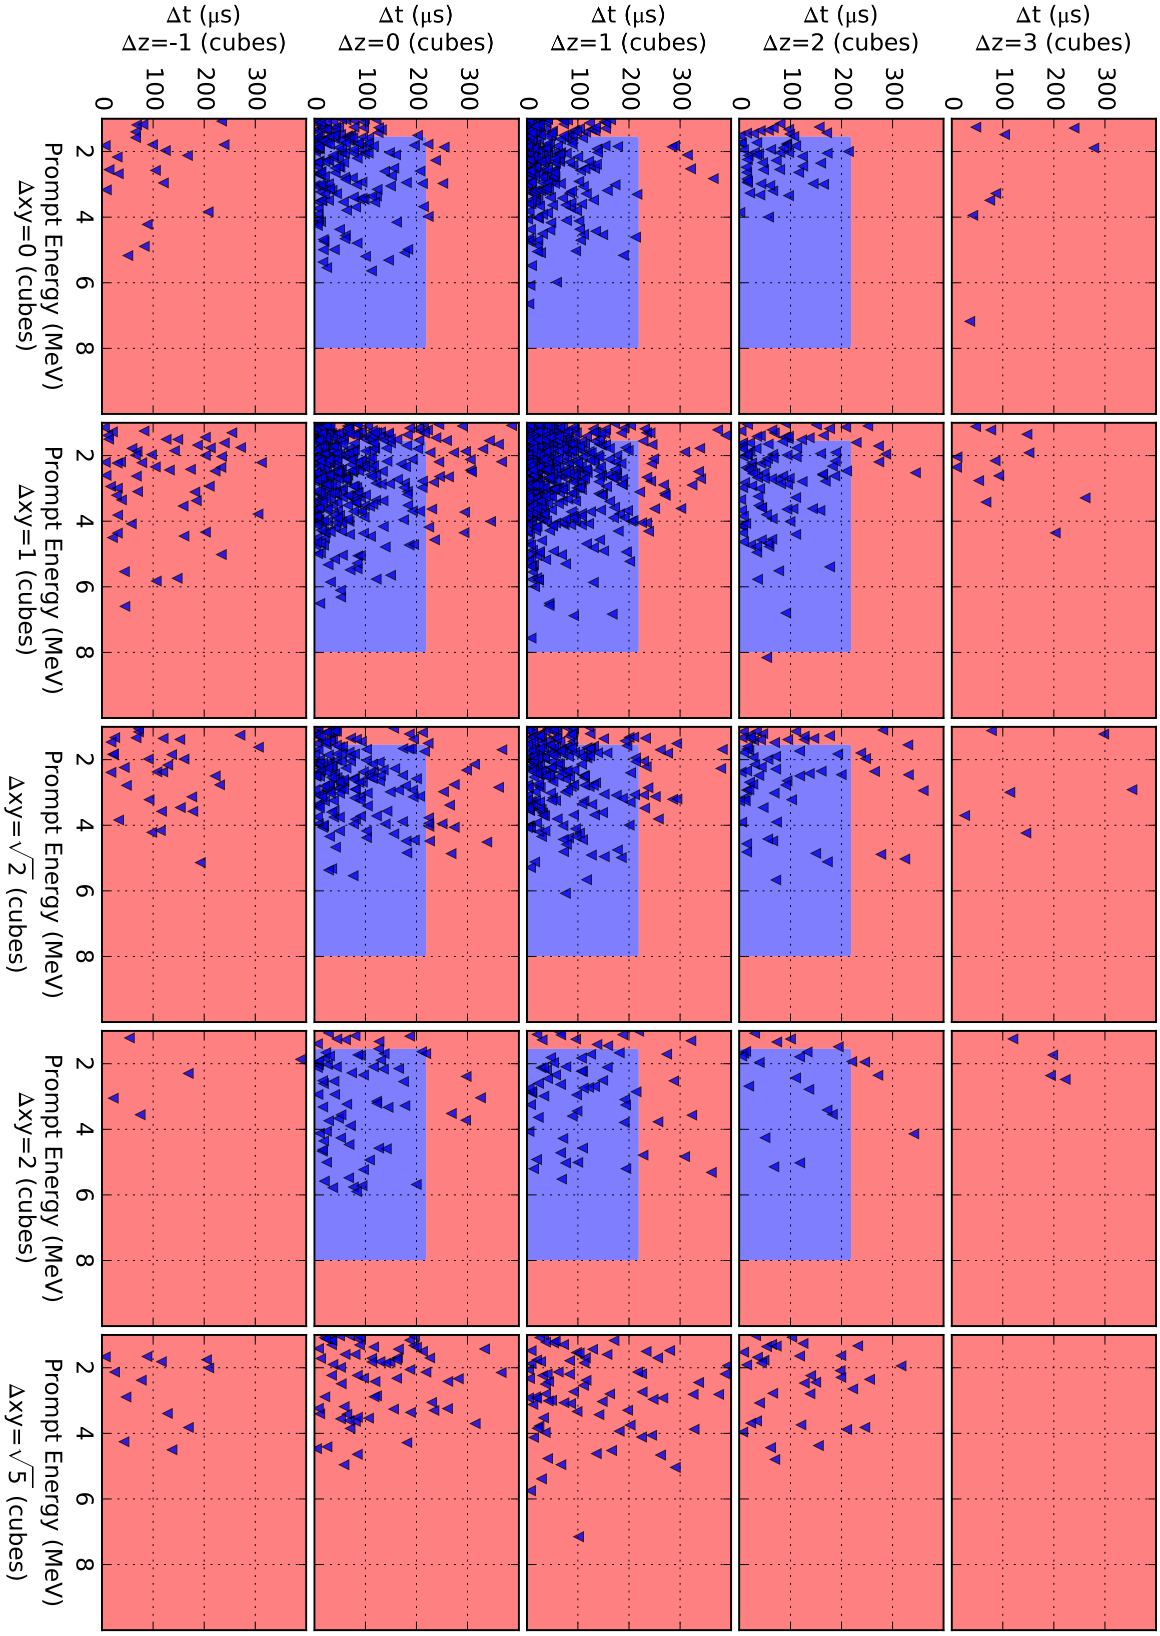

Figure 6.16: 4D distribution of IBD features for signal. The IBD selection region when using orthogonal cuts is marked in blue. The remaining region is marked in red. Note that the range plotted is to focus on signal.

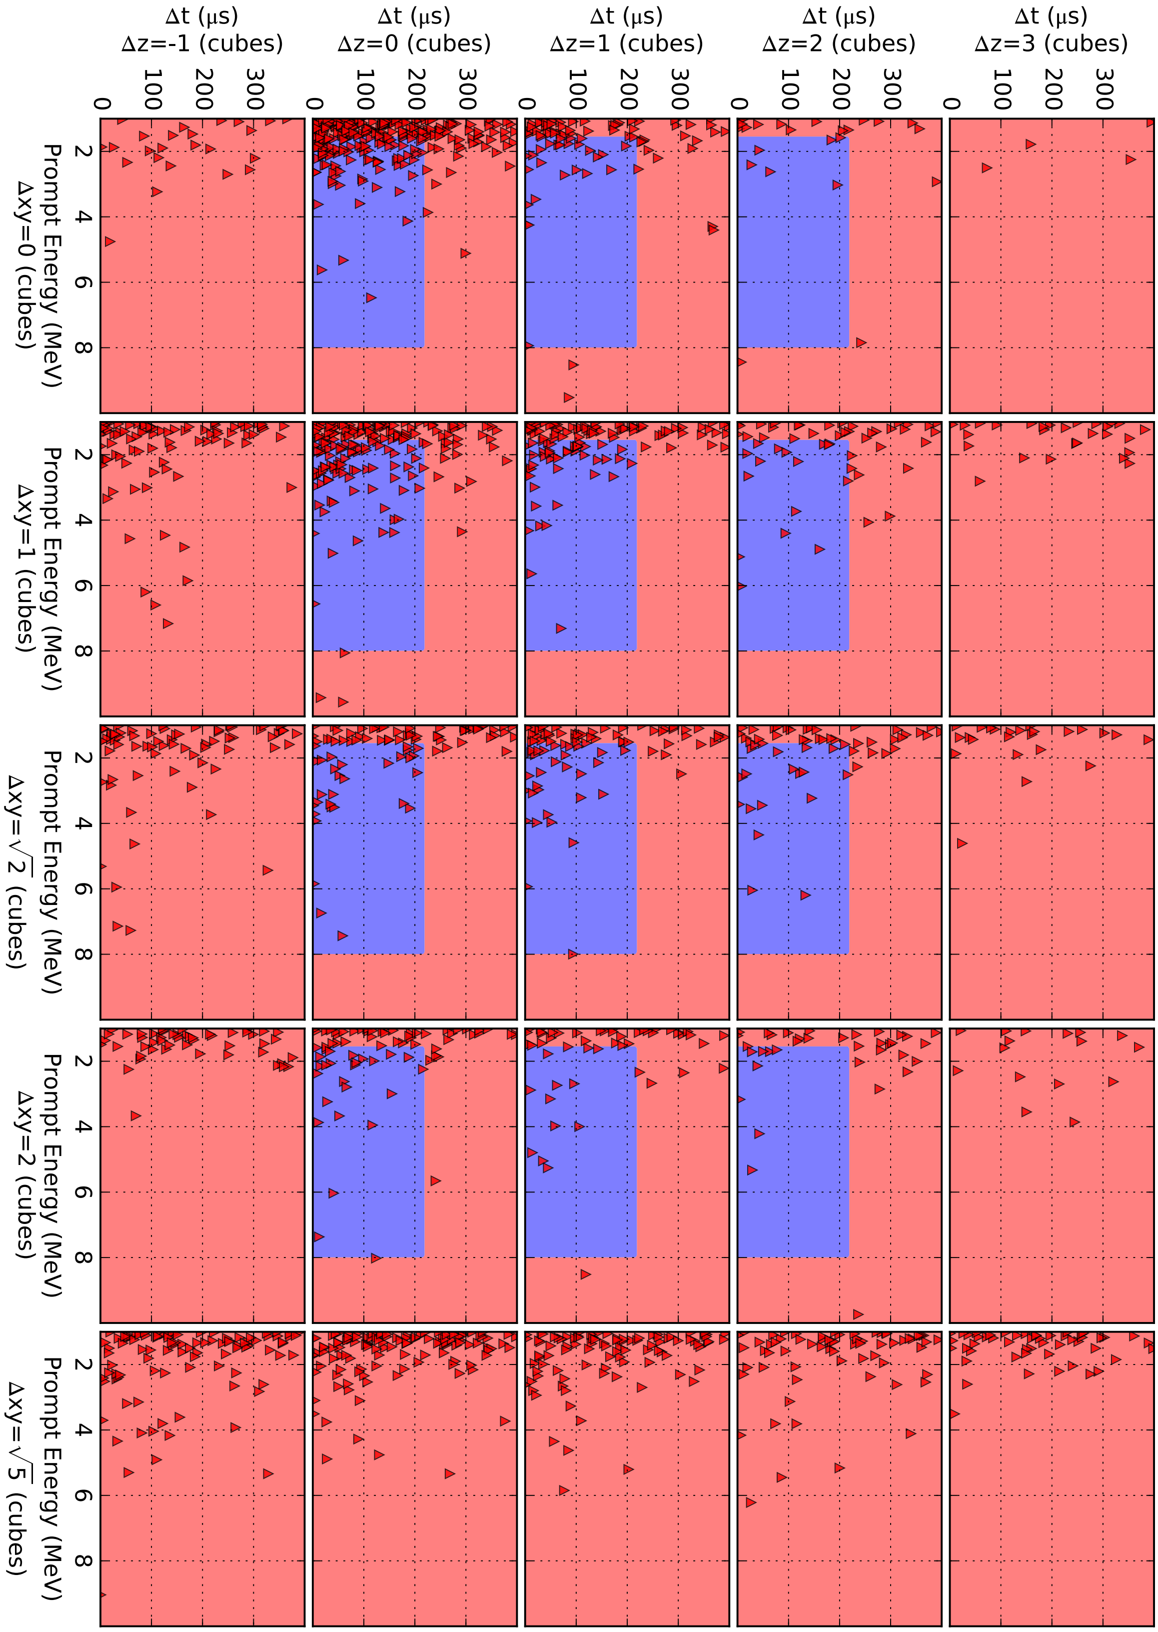

Figure 6.17: 4D distribution of IBD features for background. The IBD selection region when using orthogonal cuts is marked in blue. The remaining region is marked in red. Note that the background extends to much higher ranges than that plotted.

to classifying events. For a large dimensionality of $\mathbf{x_i}$, measuring these probability density functions from data over the whole phase space can be difficult, and become limited by the size of the dataset - this is known as the 'Curse of Dimensionality' [59]. However, many computational techniques in 'Machine Learning' (ML - AKA 'Multivariate Analysis') are available to *approximate* the likelihood function, given sets of pre-classified events. This approximated function, also known as the 'classifier', can then be used to select regions of phase space that are more likely to contain signal events compared to background, providing a more effective selection than a rectangular region.

Each machine learning technique requires some level of tuning by the user, and more sophisticated techniques that offer better separation typically require a larger investment of effort to optimise its discriminating power [59]. In this respect, two contrasting techniques are studied to find the best separation: Support Vector Machines (SVMs), and Boosted Decision Trees (BDTs). For both methods, the methodology and distinguishing features will be briefly described, and their discriminating power compared for signal classification at SoLid.

6.5.1 Training Datasets

The process of finding the approximation of the likelihood function is known as training, and the input data are referred to as the training datasets (one per classification category). Acquiring training data that is already classified can be difficult, and in practice it is common to resort to simulation. This has the disadvantage of propagating any systematic errors associated with the simulation into the training. However, reactor based neutrino experiments are very well placed to use these techniques, given the readily available large reactor off datasets that are guaranteed to be signal free. This allows the background training to be data driven. A signal training dataset is still required, and this uses the IBD simulation.

The reactor off dataset is used as the background training dataset. The classifier is therefore trained to identify the reactor off backgrounds, i.e: the sum of B_{Cor} and $B_{Acc,\,Environmental}$. The other background forming B_{On} is the accidental background from the reactor. However, there does not exist a dataset that has exclusively $B_{Acc,\,Reactor}$ candidates, and therefore this background cannot as easily be included in the background training. However, since the shape of the two accidental backgrounds is similar, a classifier trained on reactor off data will

	Dataset	N IBD Candidates	Notes
B_{Train}	Reactor Off	35K	70% Dataset available
B_{Test}	Reactor Off	15K	30% Dataset available
S_{Train}	IBD simulation	35K	N equalised to B_{Train}
S_{Test}	IBD simulation	15K	N equalised to B_{Test}

Table 6.1: The size of the test and training datasets.

still be effective at reducing this background (and any remaining accidental events studied using shifted time windows).

The signal is trained using a simulation dataset of the same size. It is useful to keep a fraction of data from entering the training, to evaluate the effectiveness of the classifier in an unbiased way. This is known as the *test* data set, and 30% of the total sample is used - the other 70% is used for the training. These fractions have been chosen to have a sufficiently large training dataset. This is summarised in table 6.1.

Pre-selection

A set of pre-selection cuts are placed on the training data sets. These are highly efficient with respect to the signal, and are used to focus the classifier on the signal region. In summary, the pre-selections are:

- $\Delta t < 400\mu$s
- $\Delta r < 5$ cubes

where $(\Delta r)^2 = (\Delta xy)^2 + (\Delta z)^2$. The signal efficiency after this pre-selection is 96%, and the background is reduced by a factor of 9.1. Unless otherwise stated, the results that follow should be interpreted relative to this pre-selection (this is also summarised clearly at the end of the section).

6.5.2 Features

It was found in section 6.4 that the four most effective IBD features for separating signal and background are: Δt, Δxy, Δz and E_{Prompt}. These are the variables used to train the following classifiers.

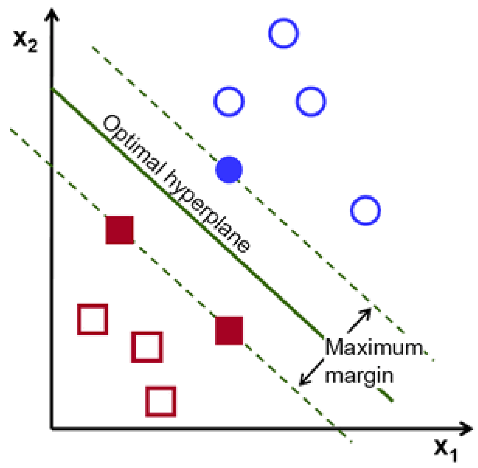

Figure 6.18: Example of an SVM decision boundary for classification of two groups (red and blue) in two dimensional feature space (x_1 and x_2). The data is arbitrary. The optimal hyperplane is found by maximising the distance between the plane and the closest data points from each category.

6.5.3 Support Vector Machines

The first method considered are support vector machines (SVMs), as described in [61]. SVM classifiers have relatively few tuning parameters, and are known to work well 'out-of-the-box' [59]. The main disadvantage of SVM classifiers are relatively slow computational performance during training, and poor scalability for larger datasets. These studies use the SVM implementation in the ML package Scikit learn [52].

Methodology

SVM classifiers work by finding linear hyperplanes that act as decision boundaries. These planes are constructed such that the distance from the nearest training data points to the planes is maximised (for both signal and background points) - see figure 6.18. In cases where data are not easily separated using linear planes in feature space, a higher dimensional space can be constructed. A kernel function is used to map between the feature space and the higher dimensional space. Many kernels are available, and depending on the shapes of the input distributions, different kernels can offer different discriminating power. A commonly used kernel is the radial bases function RBF [59], which has a single tunable parameter γ. This kernel is used for this study.

One common problem encountered during training is bias due to outliers. An additional tuning parameter C is introduced, which sets the weight attributed to each data entry. These

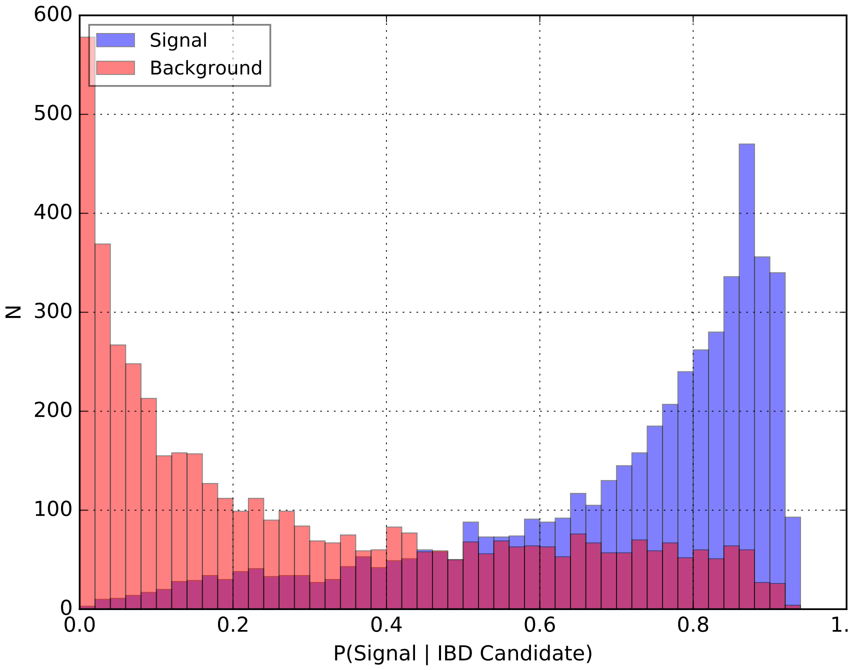

Figure 6.19: Distribution of SVM discriminating parameter for both signal and reactor off background using the standard four training variables (pre-tuning). The signal uses simulation, and the reactor off background uses data.

parameters will be tuned using data below.

The separation achieved with SVMs, with default settings, is shown in figure 6.19. This shows the distribution of the SVM discriminating variable for both signal and reactor off background. This variable is equivalent to the probability that a given IBD candidate corresponds to signal. The ROC curve is again used to judge the performance of the classifier [59], and this is shown in figure 6.20. Recall that the x axis shows the true positive rate TPR, which is the fraction of correctly tagged signal (equivalent to selection efficiency); and the y axis shows the false positive rate FPR, which is the fraction of background that is tagged as signal. The curve is drawn by measuring both FPR and TPR for different values of a cut placed on the discriminating variable. As it is desirable to have both a large TPR and a small FPR, a larger value of the ROC curve area implies better separation.

Visualisation

The classifier discriminating variable can be plotted over the IBD feature phase space, and this is shown in the 4D contour plot in figure 6.21 (using the same style as the multidimensional

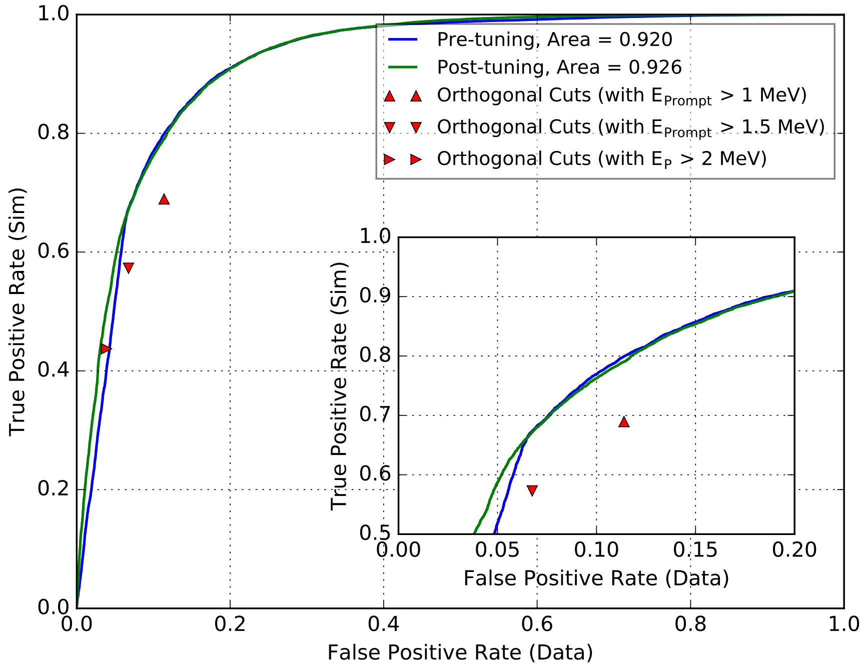

Figure 6.20: ROC curves for the 4D SVM classifier, before and after tuning. The equivalent TPR and FPR when using orthogonal cuts is also included. The upside-down red triangle shows the equivalent cut-based result from section 6.4.

visualisations from section 6.4.6). This provides many insights into the inner workings of the classifier. Given the signal candidates are mostly found in regions of high signal probability (i.e. blue points sit in blue contours), and that background points are mostly found in regions of low signal probability, the classifier is therefore behaving sensibly. The contours also show no hot spots of signal or background, and this is a good sign that the classifier has not over-fitted (AKA over-trained) the training data. Further, since ML techniques vary in how well they approximate the signal and background shapes, this visualisation provides an additional way to compare different ML techniques (e.g. by identifying particular regions in phase space where a method gives poor separation).

Figure 6.21 can be used to draw conclusions about the shape of the signal-to-background ratio in phase space. The classifier has identified regions of high signal-to-background ratio in the energy range of $2 < E_{Prompt}$ (MeV) < 6, above the large accidental backgrounds and below the correlated background at higher energies. The boost in the z direction has been identified as signal pure. Finally, at lower values of energy, the classifier identifies regions at low Δt that are likely to be signal, unlike high values of Δt, which are more likely to be accidental background. It should be noted that the colour-scale shows the decision boundaries of the classifier, and do not represent a probability density function. Blue regions show high purity, but the colour contours say nothing about the signal rate in that region. However, the scatter points plotted do give an indication of the signal and background rate in a region.

Tuning

There are two parameters that are required to be tuned to optimise this classifier:

- γ: a tuning parameter of the RGB kernel. This parameter is related to the curvature of the transformation kernel - a high value of γ results in a smoother surface.
- C: a measure of the influence of a single data point [52]. A high value of C can result in over fitting, if not compensated by a corresponding high value of γ.

These two parameters are correlated, and a scan over the ROC area over the two parameters is shown in figure 6.22. A large region of high ROC area can be seen with the ROC area at 0.93. The computational performance decreases with C, and therefore the tuning parameters are set at: $C = 20$ and $\gamma = 0.05$. The increase in ROC area are less than 1% compared to

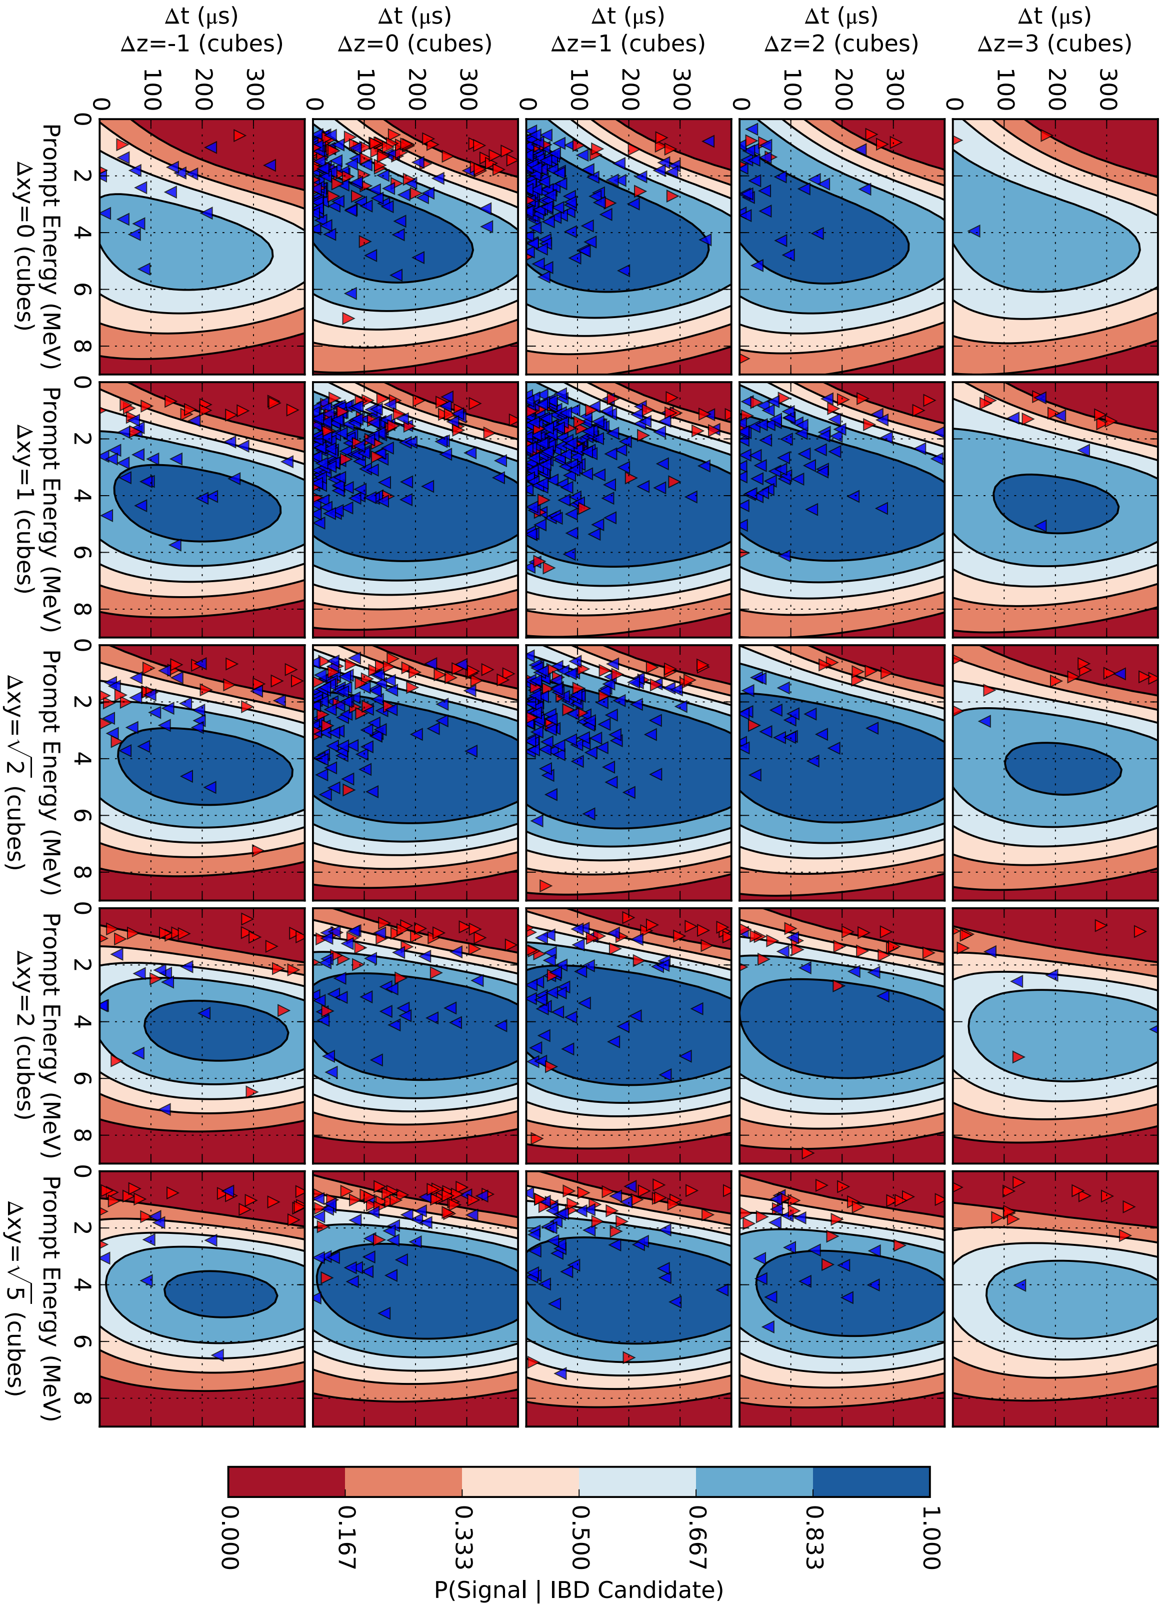

Figure 6.21: 4D distributions of signal (blue upside down triangles) and background (red upright triangle), with the SVM decision contours superposed (post tuning).

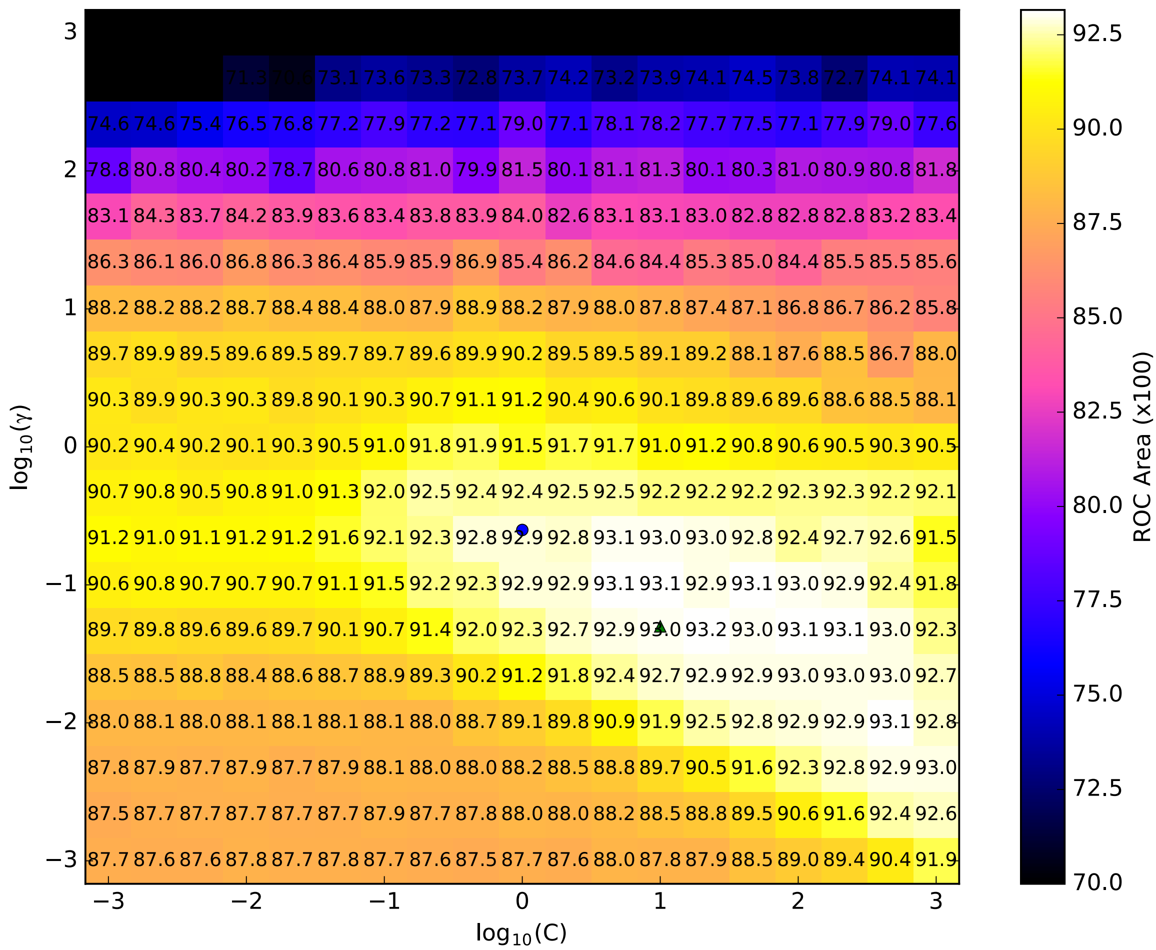

Figure 6.22: Scan of the ROC area over the two SVM tuning parameters, with the default settings marked in blue. The non-uniform colour-bar has been chosen specifically to show the features on this plot.

the default settings of the method, which is already near the maxima, suggesting the default settings of the method are close to optimal. The corresponding ROC curve for these optimal settings is also shown in figure 6.20, and is similar to the non-tuned version.

Results

Figure 6.20 also shows the equivalent FPR and TPR when using the cut-based selection from the previous section. In the region around 60% signal efficiency, the background is reduced further by $\sim 25\%$ compared to using orthogonal cuts (compare with $E_{Prompt} > 1.5$ MeV). The benefits of tuning can also be observed, where the non-tuned settings give a comparatively lower reduction of $\sim 15\%$. These reductions increase when compared to using cuts at lower energies, which is desirable for other analyses (such as the oscillation search).

One disadvantage of using ML techniques is the increased dependence on simulations to understand the signal efficiency of the classifier, and this is shown in figure 6.23 for each variable used for classification.

To understand the influence of each IBD variable in the training, the SVM classifier has been retrained by masking each of the variables in turn (this is complimentary to the cut efficiency summary plot shown in figure 6.7). This is shown in figure 6.24, for masking each IBD feature, as well as masking all topological features combined. The ROC curve with the lowest area corresponds to masking all topological features, suggesting segmentation is giving the biggest background reduction. For individual variables, the prompt energy is the most powerful discriminator. Since there are no curves sitting on top of the case where no variables are masked, we can conclude that all IBD variables are providing information that can be used to discriminate.

6.5.4 Boosted Decision Trees

Another powerful and common ML technique is boosted decision trees (BDTs). These are applied throughout experimental particle physics and beyond, and generally are one of the best methods for performing classification [52, 62]. As well as effective separation, the advantages of BDTs include computational performance and good scalability. Conversely, BDTs typically do not perform as well 'out-of-the-box', and are prone to overfitting. These studies use the AdaBoostClassifier and DecisionTreeClassifier methods, also implemented in the ML package Scikit learn [52].

Methodology

BDTs are one of several ensemble methods of ML. The goal is to combine many classifiers, known as 'weak learners', resulting in a more generalised and robust classifier. In the case of BDTs, each weak learner is a decision tree. Each learner is trained sequentially, and using the full training dataset. Decision trees themselves are one of the simplest ML techniques. However, they can be prone to over-fitting, as well as becoming overly complex for simple datasets. As a result, individual trees by themselves are not especially effective, but these effects can be overcome by using them in an ensemble.

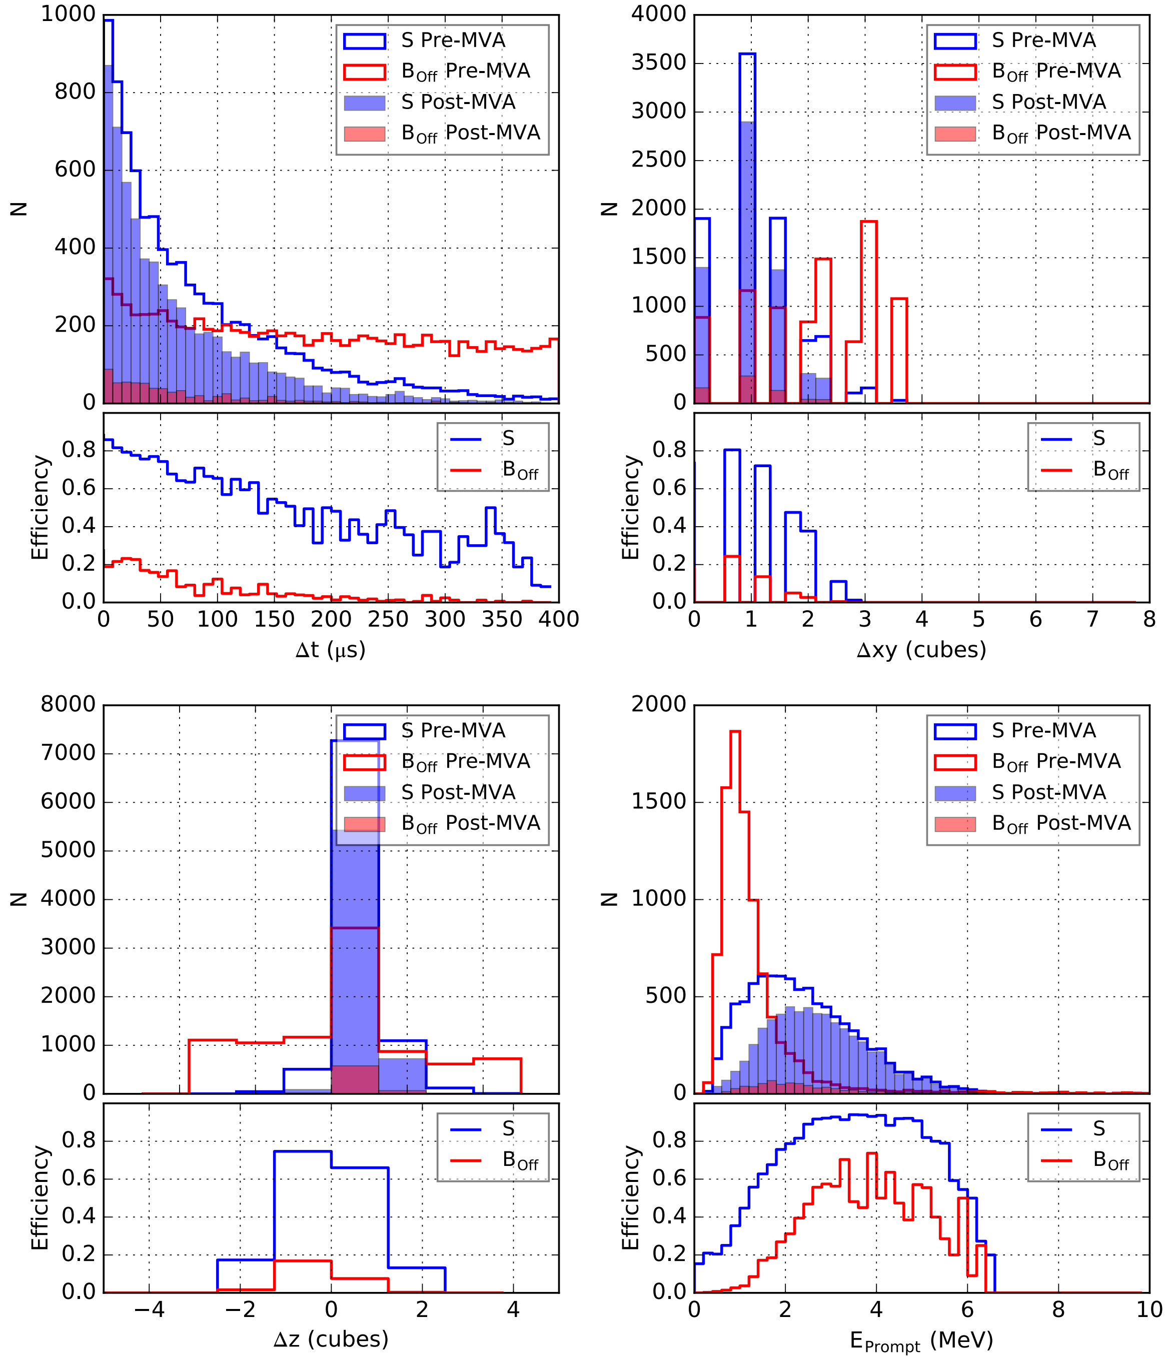

Figure 6.23: Selection efficiency for each IBD variable used for classification for the SVM classifier.

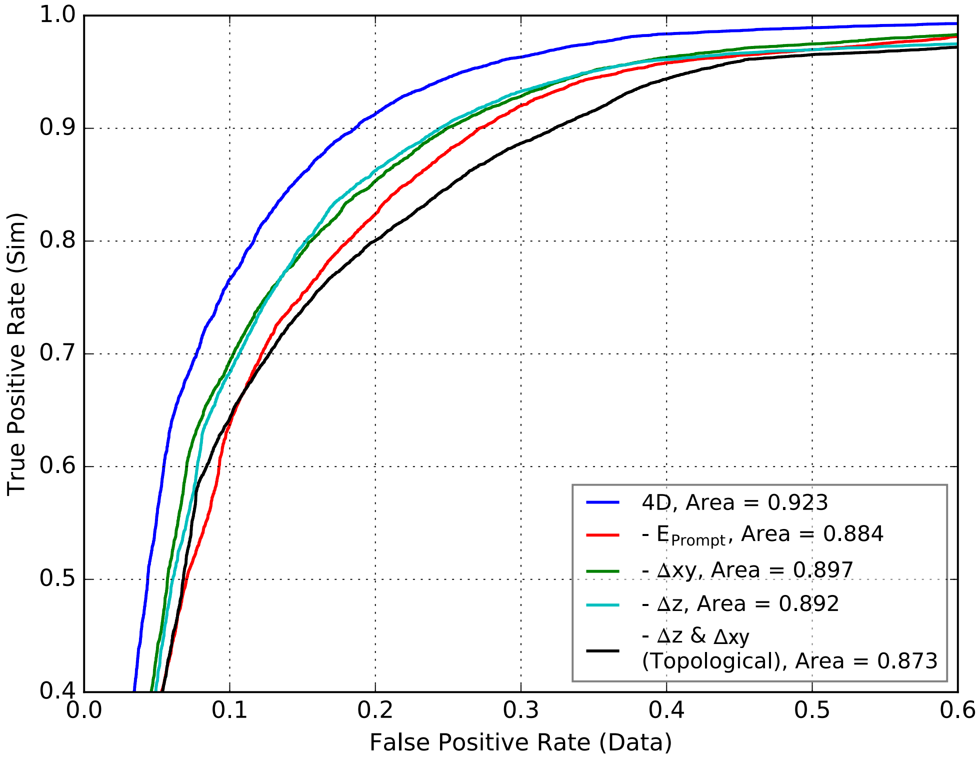

Figure 6.24: ROC curves using SVMs when masking particular variables during the training of the classifier, for signal (simulation) and background (reactor off data).

At the beginning of the training of the BDT, each data entry of the training dataset is assigned a weight. During the training of the first weak learner, all weights are equal. These weights are modified during the training of the subsequent learners, such that entries found in regions of phase space that are fitted wrongly are given more emphasis. Consequently, each learner in the sequence is forced to better fit the poorly fitted regions of the previous learner, possibly at the expense of becoming less effective in other regions. As a result, particular learners are especially effective for particular regions of phase space. The output of all weak learners is combined via a final weighted sum when making classifications.

Optimisation

There are two main parameters that require to be tuned [52]:

- **Number of Estimators**: the number of weak learners.
- **Maximum Decision Tree Depth**: the upper limit on the depth of each decision tree learner. This parameter controls the level of over fitting.

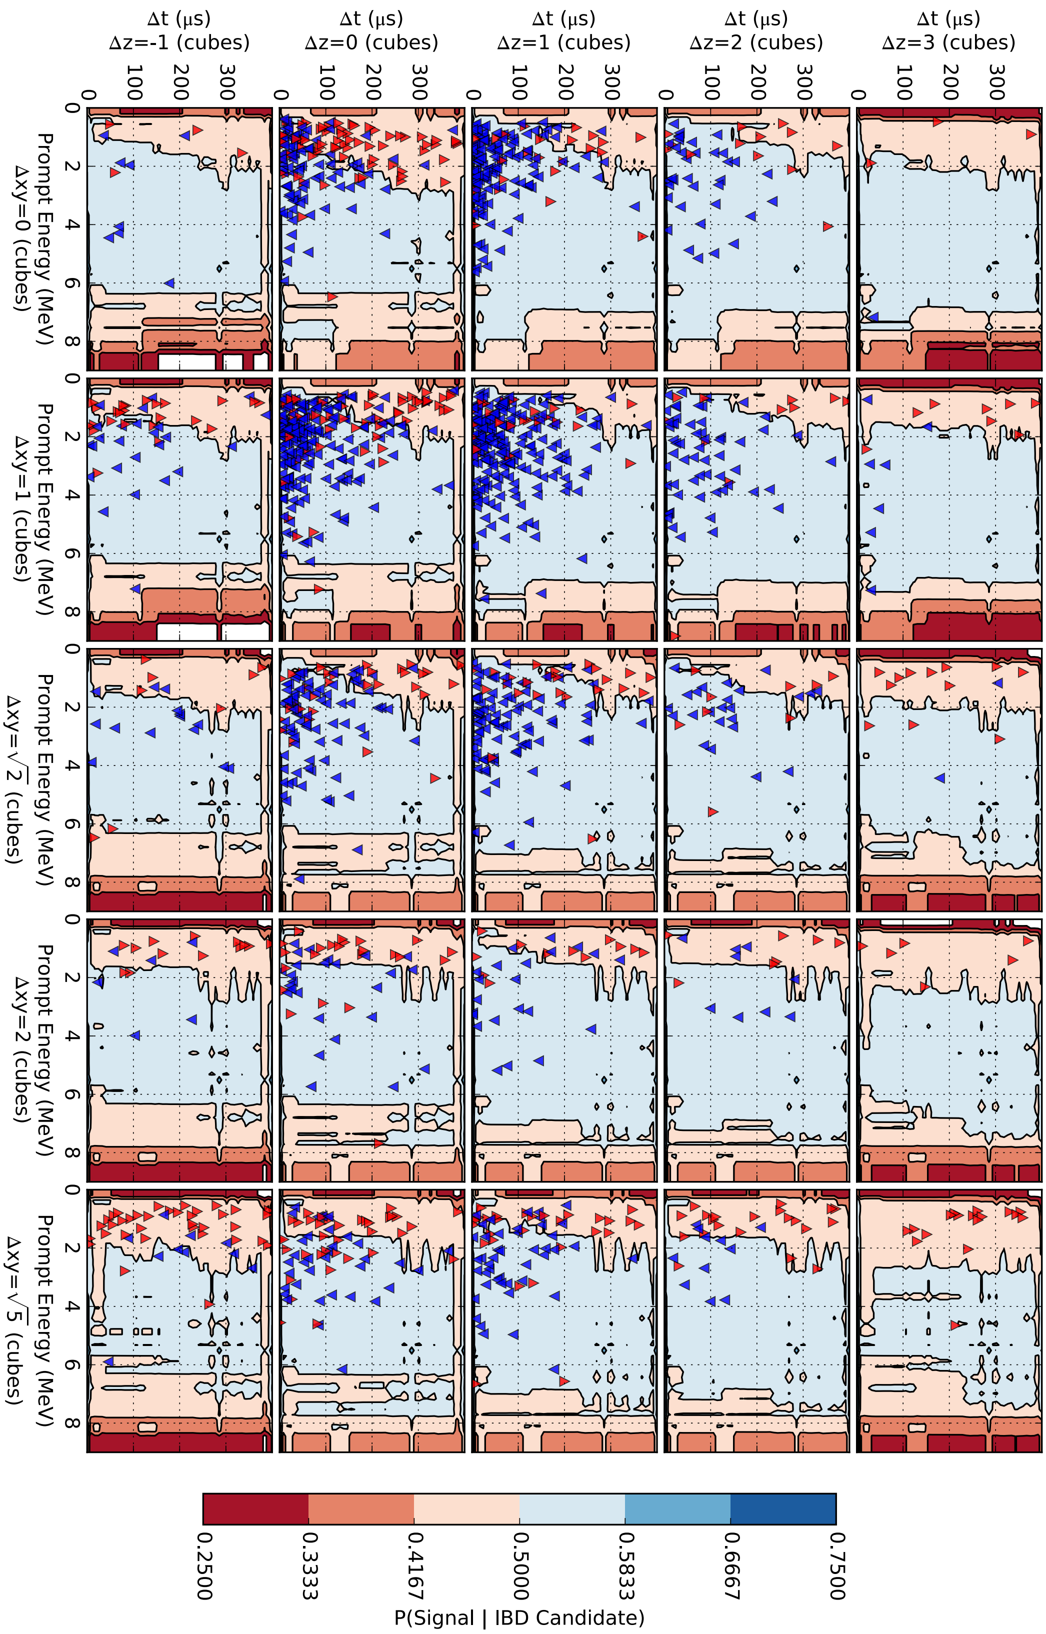

Figure 6.25: 4D distributions of signal (blue upside down triangles) and background (red upright triangle), with the BDT decision contours superposed (post tuning). The 'boxy' shapes are expected since the BDT is an ensemble of decision trees - see also [52] for similar examples.

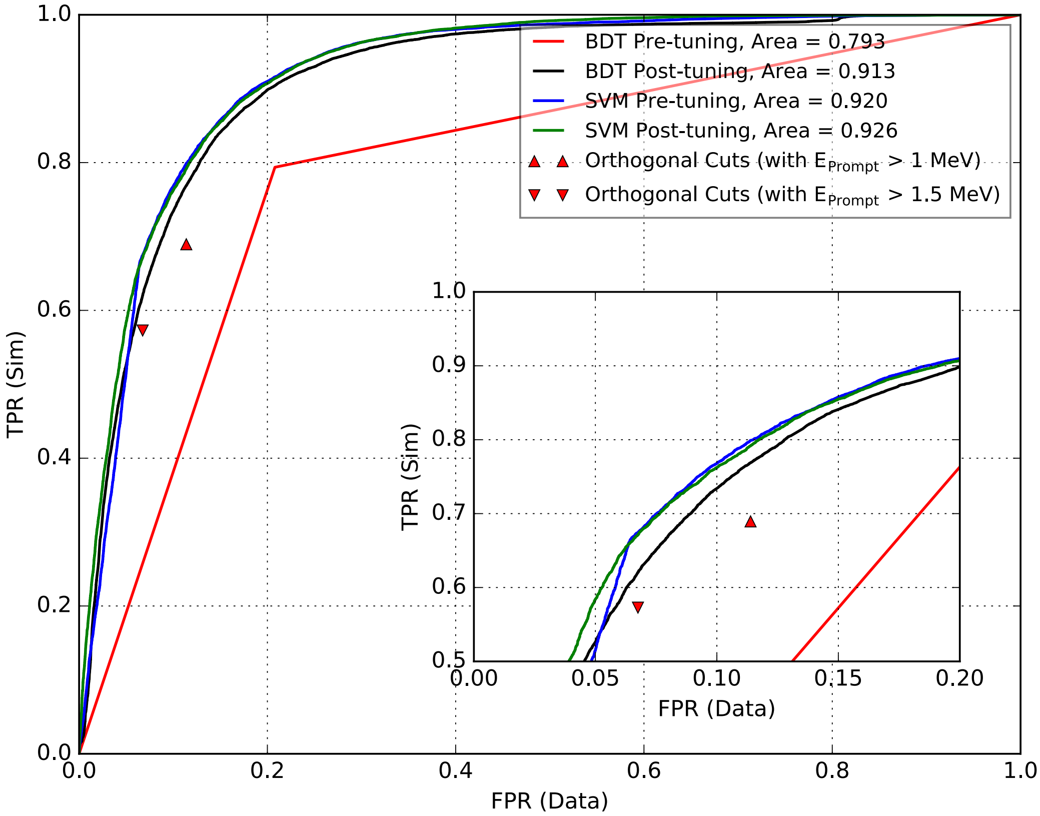

Figure 6.26: ROC curves for both BTDs and SVMs. The low area and jagged curve for the non-tuned BDT is caused by both the low number of the weak learners and maximum depth when using default settings.

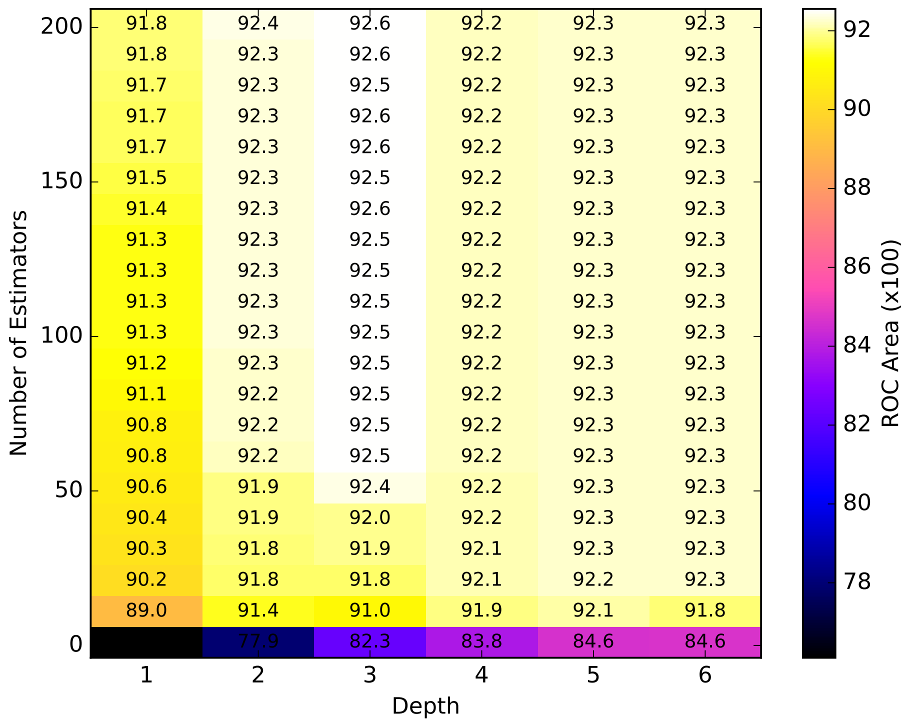

Figure 6.27: Optimisation of the BDT classified. A scan of the ROC curves area over the two main BDT tuning parameters.

A scan of ROC area over these two parameters is shown in figure 6.27. The maximum is at higher values of the number of estimators, and at a maximum depth of 3. The CPU time increases with the number of estimators, and therefore the tuned parameters are set as: maximum depth = 3, and number of estimators = 150. It can also be seen the variation in ROC area is greater than that found for SVMs. A 4D visualisation of the BDT decision contours is shown in figure 6.25, which shows similar general patterns, although differs in the shapes used to describe these patters (e.g more 'boxy' compared to the smoother contours for the SVM classifier).

The ROC curves corresponding to these tuned parameters, and the default 'out-of-the-box' settings for the BDT are shown in figure 6.26. The tuned BDT classifier performs significantly better once tuned. Figure 6.26 also shows the results obtained with the SVM classifier from above. Since the tuned performance of the BDT classifier is less optimal than that of the tuned SVM classifier, the tuned SVM classifier is the classifier used for the reactor on/off comparison described in section 6.6.

6.5.5 Summary

Two machine learning techniques, SVMs and BDTs, have been investigated and tuned to separate IBD signal from reactor off backgrounds. The most effective method, based on ROC area measurements, are tuned SVMs. This classifier can reduce around 25% more background when compared to the cut-based selections. The most powerful single discriminating variable used in training is E_{Prompt}. The topology of the event gives the best discriminating information, as can be seen when considering all topological variables combined. The performance of SVMs depends marginally on tuning, unlike BDTs, whose performance improves considerably post tuning. Multidimensional visualisations have demonstrated that these classifiers are performing sensibly, and that they identify many of the expected trends for the signal-to-background ratio.

6.6 Reactor On Results

Both the cut-based selection, and the tuned SVM classifier, are now applied to the reactor on data set to search for a neutrino signal. The following two hypotheses will be compared:

- **Null hypothesis**: there are no IBD candidates from reactor neutrinos.
- **Alternative hypothesis**: there are IBD candidates from reactor neutrinos.

There are multiple tests that can be used to compare these two hypothesis - two are considered:

- **Rate**: the number of IBD candidates found in the reactor on dataset is compared with the background expectation. A neutrino signal would be observed as an excess in the number of correlated events. This is performed for both the cut-based and the SVM selections separately.
- **Rate + Shape**: as well as using the rate information, the shape of the IBD feature distributions can also be utilised. An example test using both types of information is the 'extended likelihood ratio' test [66]. This is an extension of the standard likelihood ratio test, which also incorporates the rate of the hypothesis, thus using the rate in addition to the shape information.

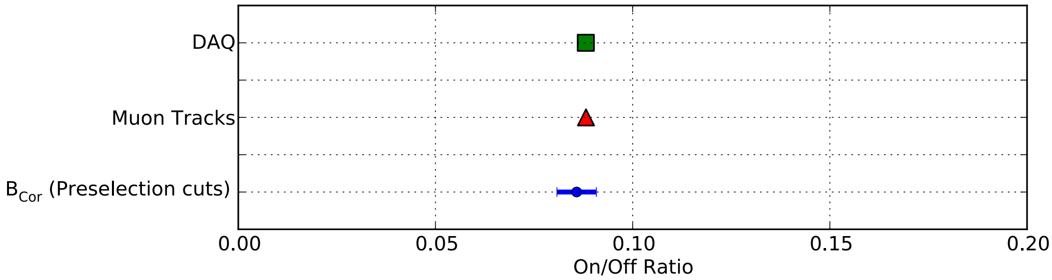

Figure 6.28: Comparison of methods used to measure the ratio of reactor on to reactor off detector live time. For the correlated background, only the pre-selection cuts were applied, dominating any correlated signal in the reactor on data.

The power of these tests, and the significance of any observed signal (or lack thereof) is discussed at the end of this section.

6.6.1 Background Model

Recall that the background model for the reactor on period B_{On} is the sum of:

- $B_{Acc,\,On}$: found by using shifted time windows during the IBD reconstruction of reactor on data.
- B_{Cor}: found using the reactor off dataset, with the environmental accidental background subtracted ($B_{Cor} = B_{Off} - B_{Acc,\,Environment}$, with $B_{Acc,\,Environment} = B_{Acc,\,Off}$).

The background rates found from the reactor off dataset need to be scaled according to the ratio of the live times of the reactor on and off runs. This value is provided by the DAQ of the prototype, and is found to be 0.082. To cross check this result, this ratio can be compared to other ratios that are expected to be independent of the reactor power, such as the number of crossing muons, or the correlated background rate (using pre-selection cuts only). Their ratios are shown in figure 6.28, and are found to be consistent with that supplied by the DAQ.

6.6.2 Signal Prediction

The predicted number of IBD interactions $\langle S \rangle$ that would pass the above selection criteria and trigger is taken as:

$$\langle S \rangle = N_{IBD\ Interactions} \times \epsilon_{Analysis} \times \epsilon_{Detection} \tag{6.6}$$

where $N_{IBD\ Interactions}$ is taken as the prediction from the reactor calculations (equation 6.1), and $\epsilon_{Analysis}$ is the efficiency of the analysis selection criteria. For the cut-based selection, this is 57%, and the signal acceptance of the SVM classifier is set to match this efficiency. $\epsilon_{Detection}$ is the efficiency to detect and trigger IBD candidates. Since the prototype has high efficiency for detecting prompt events, this is dominated by the neutron detection efficiency: $\epsilon_{Detection} \approx \epsilon_{Neutron}$. This is the product of the neutron capture efficiency, and trigger efficiency. This has been been previously measured using the AmBe dataset [49, 55], and it is found $\epsilon_{Neutron} = 2.8\%$. Therefore:

$$\langle S \rangle = 10.0 \text{ Interactions} \tag{6.7}$$

The systematic uncertainties in this calculation are discussed later in section 6.6.5.

6.6.3 Cut-Based Selection

Rate

A summary of the background expectation rates, and the number of observed IBD events for the reactor on period using the cut-based selection criteria, is shown in table 6.2. A small excess of 10 ± 17 (stat) IBD interactions is observed. This error is dominated by the uncertainty of the background rate, which is Poisson distributed[3]. The excess is therefore equivalent to a 0.6σ upward fluctuation of the background, and is not statistically significant.

The corresponding distribution of Δt is shown in figure 6.29, which is consistent with this finding. This figure also shows how the backgrounds have been modelled appropriately.

Rate + Shape

The prompt energy distribution is shown in figure 6.30. The data show no significant outliers compared to the background expectation by eye, suggesting again that the background model

[3]Note that this is not limited by the size of the datasets used to predict the background rate - the uncertainty of the expected background rate is small relative to the fluctuations in this rate due to the fact the rate is Poisson distributed.

Component	N IBD Candidates (scaled to Reactor On live time)		Source
	Cuts	SVM	
B_{Off}	258 ± 5	216 ± 4	Measured
$B_{Acc,\,Environmental}$	25 ± 1	27 ± 1	Measured using shifted time windows (reactor off)
B_{Cor}	234 ± 5	189 ± 4	B_{Off} - $B_{Acc,\,Environmental}$
$B_{Acc,\,On}$	61 ± 2	58 ± 3	Measured using shifted time windows (reactor on)
$B_{Acc,\,Reactor}$	37 ± 2	31 ± 3	$B_{Acc,\,On}$ - $B_{Acc,\,Environmental}$
$\langle S \rangle$	10 ± 1	10 ± 1	Reactor calculations
$\langle B_{On} \rangle$	295 ± 5	247 ± 5	$B_{Acc,\,On} + B_{Cor}$
$\langle S : N \rangle$	1:30	1:25	-
$N_{Reactor\ On}$	305	263	-
Excess	10 ± 18	16 ± 17	-
Significance	0.6σ	0.9σ	-

Table 6.2: Background expectation and measurements of the IBD candidate rate for the reactor on period. Errors are statistical only (assuming Poisson statistics [58]). Coincidently, the excess is very near the expected excess, but given the errors, this is not statistically significant.

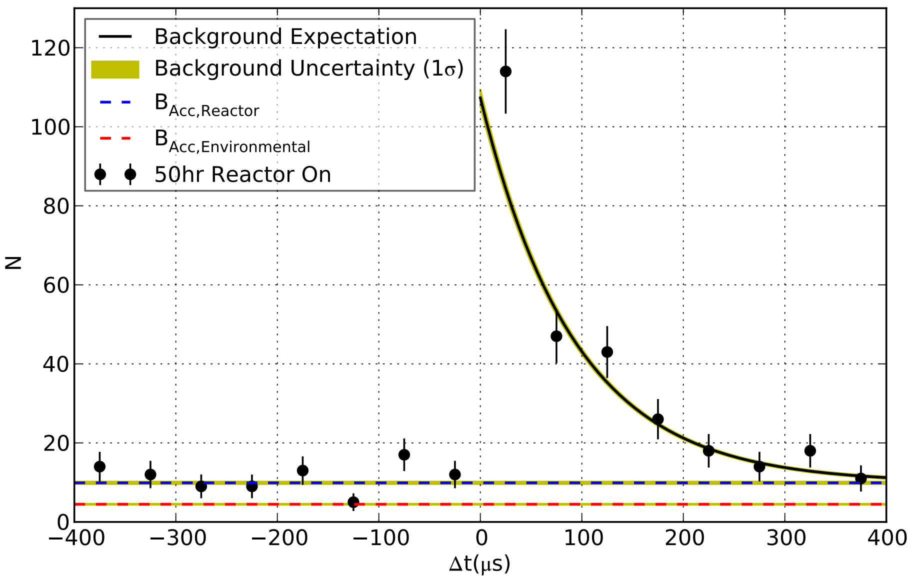

Figure 6.29: Δt distribution for the reactor on dataset (post cut-based selections). The shape of the neutron capture time (for the correlated background) has been fitted previously using a single exponential, as outlined in [49] (included for aesthetics only). The background histograms are stacked.

has been predicted accurately. An extended likelihood ratio test is performed using this distribution.

In the general case, the extended likelihood parameter $\mathcal{L}_{Ex}$ takes the form [66]:

$$\mathcal{L}_{Ex} = \prod_{i=1}^{N} p(\mathbf{x_i}) e^{-\langle N \rangle} \frac{\langle N \rangle^N}{N!} \tag{6.8}$$

In this analysis, N corresponds to the number of IBD candidates observed, and $\langle N \rangle$ is the expectation. Since the expectation depends on the amplitude of the neutrino signal A: $\langle N \rangle \to \langle N \rangle(A)$. A is allowed to float. In the general case, $p(\mathbf{x_i})$ is the normalised probability distribution function, sampled for event $\mathbf{x_i}$. It can be seen this equation differs from the standard likelihood parameter by the additional Poisson probability term. In this analysis, $p(\mathbf{x_i})$ is the prompt energy distribution of IBD candidates from figure 6.30, with the signal contribution scaled by the parameter A. Therefore, $\mathcal{L}_{Ex}$ becomes:

$$\mathcal{L}_{Ex} = \prod_{i=1}^{N} p(E_{Prompt,i}; A) e^{-\langle N \rangle (A)} \frac{\langle N \rangle (A)^N}{N!} \tag{6.9}$$

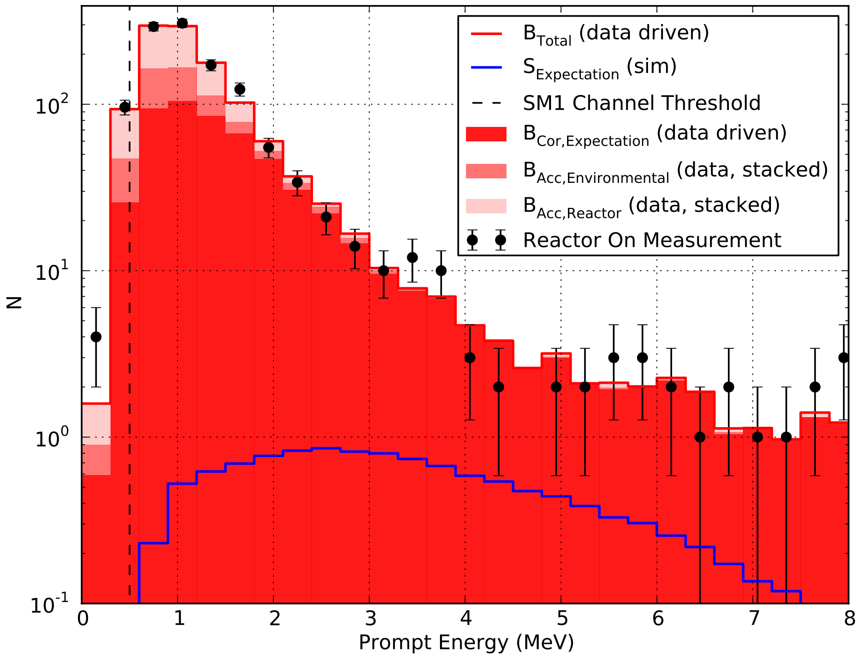

Figure 6.30: Prompt energy distribution for the reactor on dataset (post cut-based selections). Curves are stacked. Events found below the average trigger threshold are present due to variations in the online equalisation of the detector channels.

It is computationally beneficial to use the log of the likelihood (to avoid many products of small terms). This equation can then be re-arranged to give [47]:

$$\ln(\mathcal{L}_{Ex}) = N\ln(\langle N\rangle(A)) - \langle N\rangle(A) + \sum_{i=1}^{N} \ln p(E_{Prompt,i}; A) \tag{6.10}$$

To find the best fitting value of A, a scan has been performed and is shown in figure 6.31. The scan is over the likelihood ratio: $2\ln(\mathcal{L}_{Ex,S+B} - \mathcal{L}_{Ex,B})$, and compares the two hypotheses being tested (signal plus background vs. background only). The maximum is at $A = -8 \pm 13$ (stat) IBD interactions for the reactor on period, where the statistical uncertainty is measured by searching for a change in the log-likelihood of 1. The significance of this value is discussed in section 6.6.5.

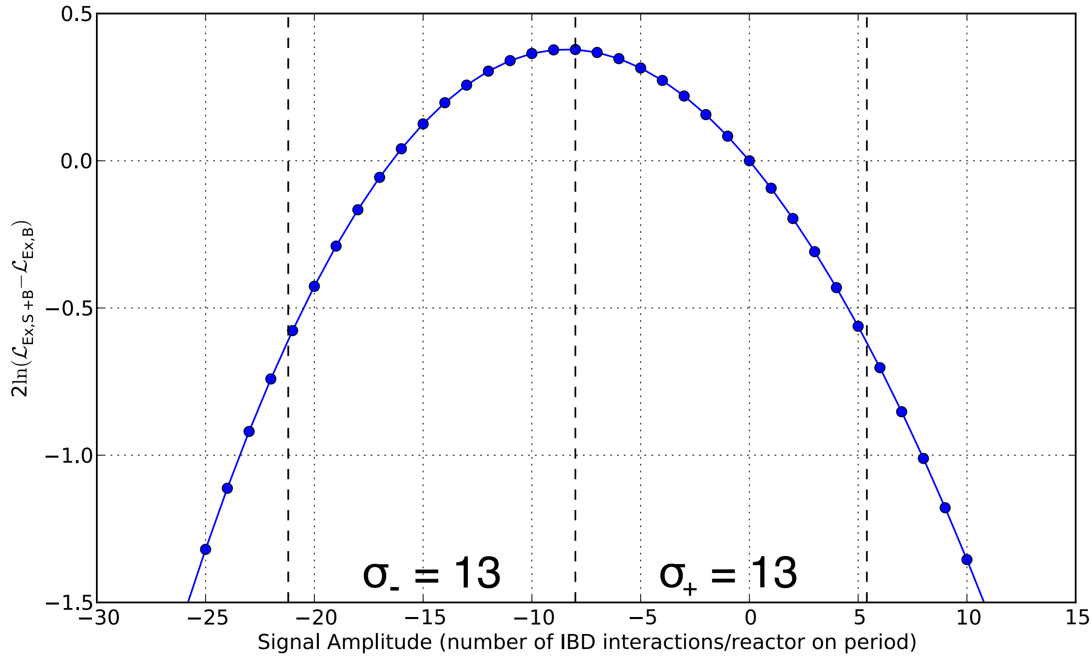

Figure 6.31: Results of the rate plus shape analysis. A scan of the extended likelihood ratio over signal amplitude, to find the best fitting value.

6.6.4 Support Vector Machines

Rate

Table 6.2 also shows the signal and background rates for the reactor on period when using the SVM classifier. Since the output of the discriminator is a continuous variable, a rejection threshold is required to be set to find these rates. To compare with the previous cut-based selection, the threshold is set at 0.84 to have equal $\epsilon_{Analysis}$ to that of the cut-based selection. This is also marked in figure 6.32, which shows good agreement between the data and background model, with few outliers in the region of high signal probability (by eye).

The rates of the accidental backgrounds are consistent within error with the rates found using the cut-based selection, and the correlated rate is decreased. The overall background is reduced by approximately 20% compared to the cut-based selection. Similarly, a small excess of 0.9σ is observed for the reactor on period.

6.6.5 Discussion

Systematic Errors

The following three sources of systematic error are explored:

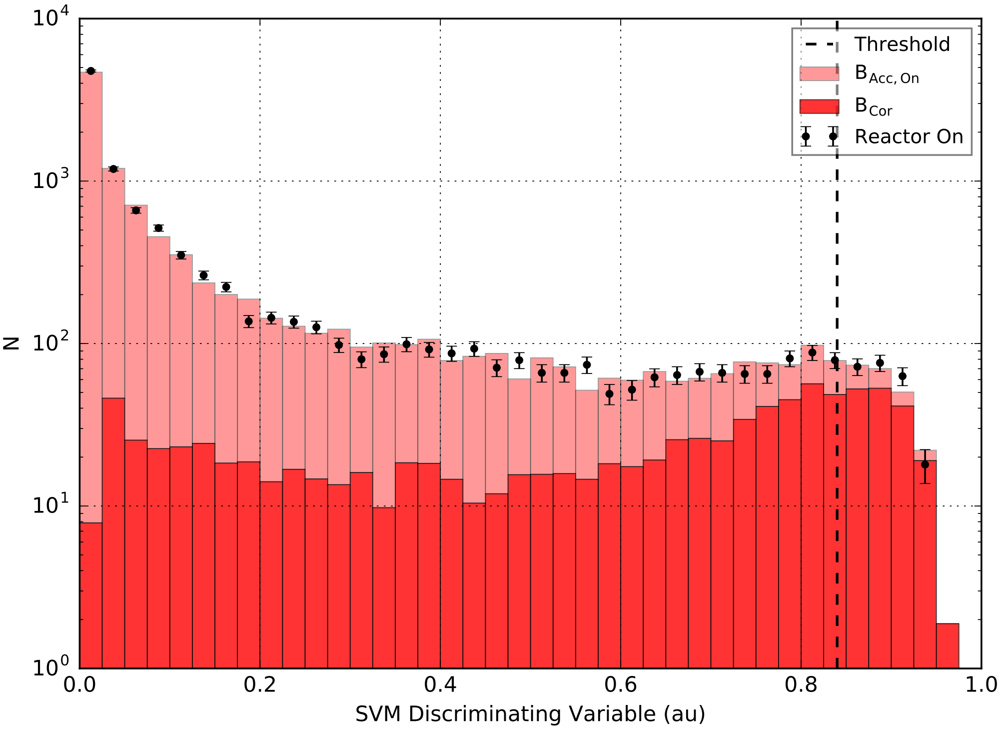

Figure 6.32: SVM discriminating variable for the reactor on dataset. Entries are used for the reactor on analysis signal selection if above 0.84.

- **Reactor power**: the dominant error in predicting the absolute neutrino rate emitted by the reactor arises from the reactor power measurements. Dedicated sensors placed at the reactor, which measure the energy flow of the system, are accurate to $\sim 8\%$ [67]. These measurement techniques have been designed for power reactors, where greater precision is typically not required. For BR2, this precision is expected to improve in future with improved calibration data for these sensors.

- **Neutron efficiency**: the calculation of the neutron efficiency is accurate to $\sim 15\%$ [49]. This measurement has been made using an AmBe source placed at multiple positions around the prototype. There are two main contributions to this uncertainty:

  - The calculation of detector dead-time, which has increased fluctuations compared to reactor off data due to the high rate of the AmBe source itself.
  - The AmBe source activity, which is known to $\sim 10\%$ accuracy.

- **Readout simulation energy scale**: the parameters of the readout simulation have been tuned to data. The simulation of IBD interactions is dependent on the accuracy of these parameters. In particular, the absolute energy scale of the simulation is accurate to the $\sim 20\%$ level [68]. To propagate this uncertainty into the IBD analysis, the number of predicted IBD candidates has been recalculated by varying the value of the prompt energy cut by $\pm 20\%$ - see table 6.3. It can been seen that the change in the signal prediction is 1.2 events (in both directions), and so this is taken as the uncertainty on the number of predicted IBD candidates: $\sim 12\%$ error.

Taking the combined systematic uncertainty to be the quadrature sum of these contributions, the number of predicted IBD interactions during the reactor on period is assigned the following systematic uncertainty:

$$\langle S \rangle = 10 \pm 3 \text{ (syst)} \tag{6.11}$$

It can be seen that the value of the systematic uncertainty is significantly lower than the statistical uncertainty, which is dominated by the subtraction of the large background.

E_{Prompt} Lower Cut (MeV)	Cut Efficiency (%)	$\langle S \rangle$	Remark
1.5	74	10.0	Default case
1.2	83	11.2	Sim energy scale under estimated by 20%
1.8	64	8.8	Sim energy scale over estimated by 20%

Table 6.3: The number of IBD interactions for the reactor on period $\langle S \rangle$ for a variety of lower prompt energy cuts. This emulates the case where the errors in the readout simulation cause the energy scale to be under-estimated or over-estimated.

Statistical Significance

To assess the statistical significance of the three tests conducted above, the experiment was simulated many times, for both the background only hypothesis B and the background plus signal hypothesis $B + S$, to draw distributions of the test statistics. For the rate tests, a random signal and background rate is generated based on the expectations outlined in table 6.2. For the shape test, the signal and background events are randomly generated using the prompt energy distribution (with the signal amplitude fixed to the prediction).

Figure 6.33 (left) shows the distributions of the rate and SVM test statistics for a 2-day reactor on run. The results found above are also plotted as vertical lines. There is significant overlap between the two hypotheses, for both the rate and SVM test statistic. From this, we can conclude that the data shows no preference for either hypothesis - i.e. the data is compatible with both hypothesis, and the dataset ins insufficient to confidently classify one of the two hypotheses. Figure 6.35 shows a similar situation for the rate plus shape test using the extended likelihood ratio test.

The right of figure 6.33 shows the case for a 20-day run, and these distributions are shown as a function of time in figure 6.34. It can be seen that after a 20-day run, the prototype detector would most likely observe a neutrino signal to 2σ confidence level, and this increases with time.

On the power of the tests used, the separation between the two hypotheses is slightly increased for the extended likelihood test compared to the other two tests. Since this is a classification problem, this can again be demonstrated using ROC curves, as shown in figure 6.36 for all three test statistics. The area under the curve for the extended likelihood ratio test is highest

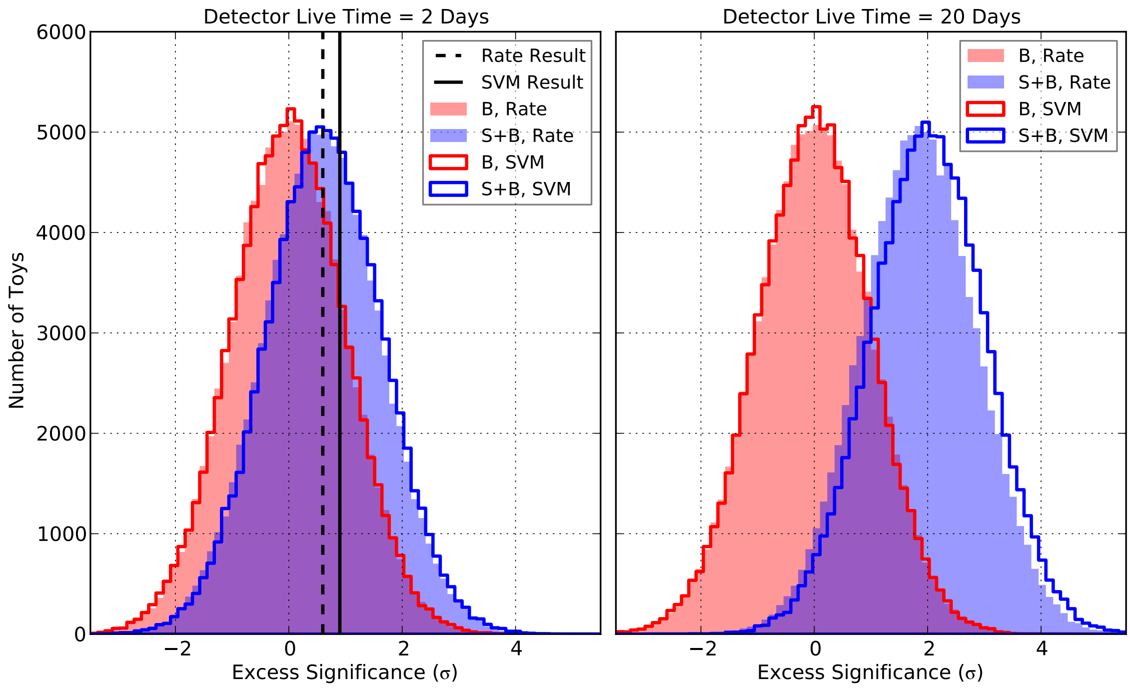

Figure 6.33: Distributions of the different tests statistics used to search for the neutrino signal, from simulation of the prototype reactor on run. The shaded distributions correspond to the rate test using the cut selection, and the clear distributions (only the bins outlined) show the rate test using the SVM classifier. Left: a 2-day run, equivalent to the data available. The results from data are shown as black vertical lines. Right: the case had the prototype detector ran for 20 days with the reactor on.

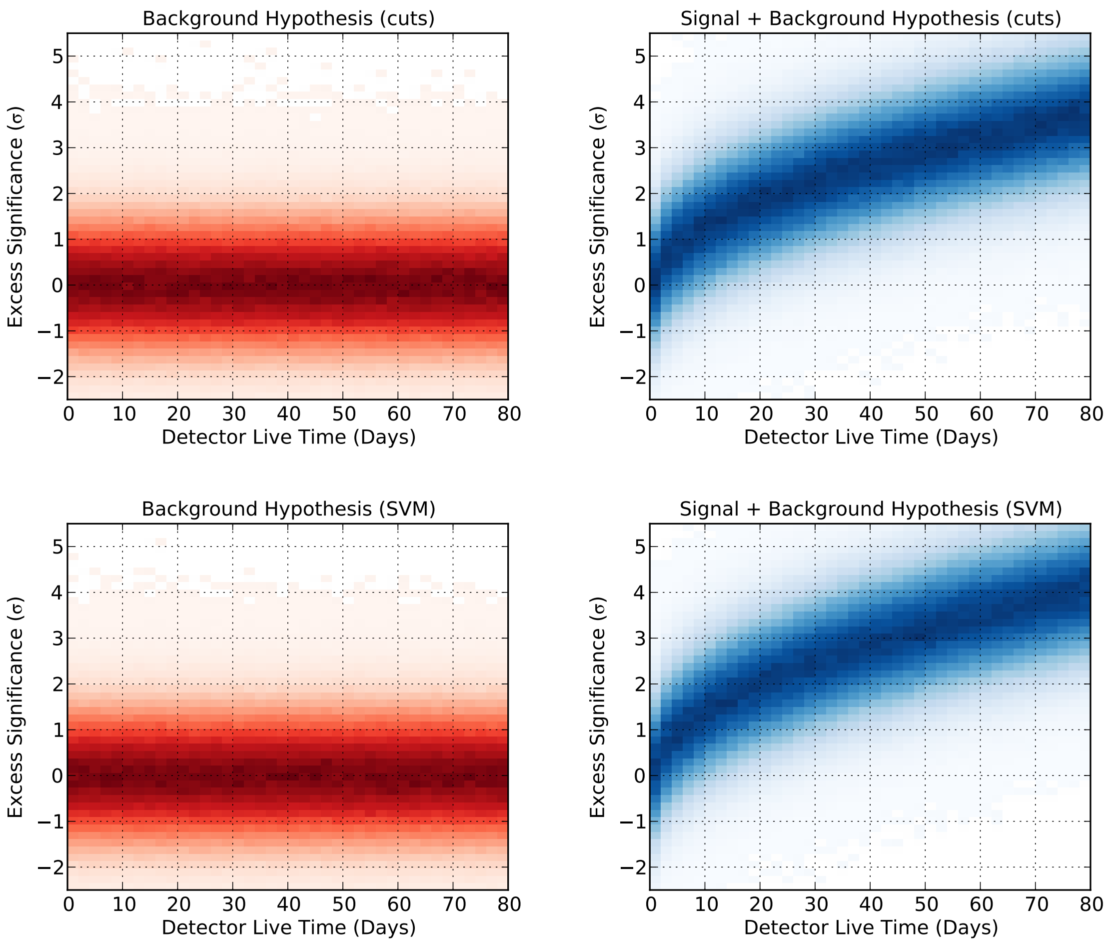

Figure 6.34: Simulated distributions of the different tests statistics as a function of detector live time (reactor on). Slices at $t = 2$ days and $t = 20$ days are shown in figure 6.33.

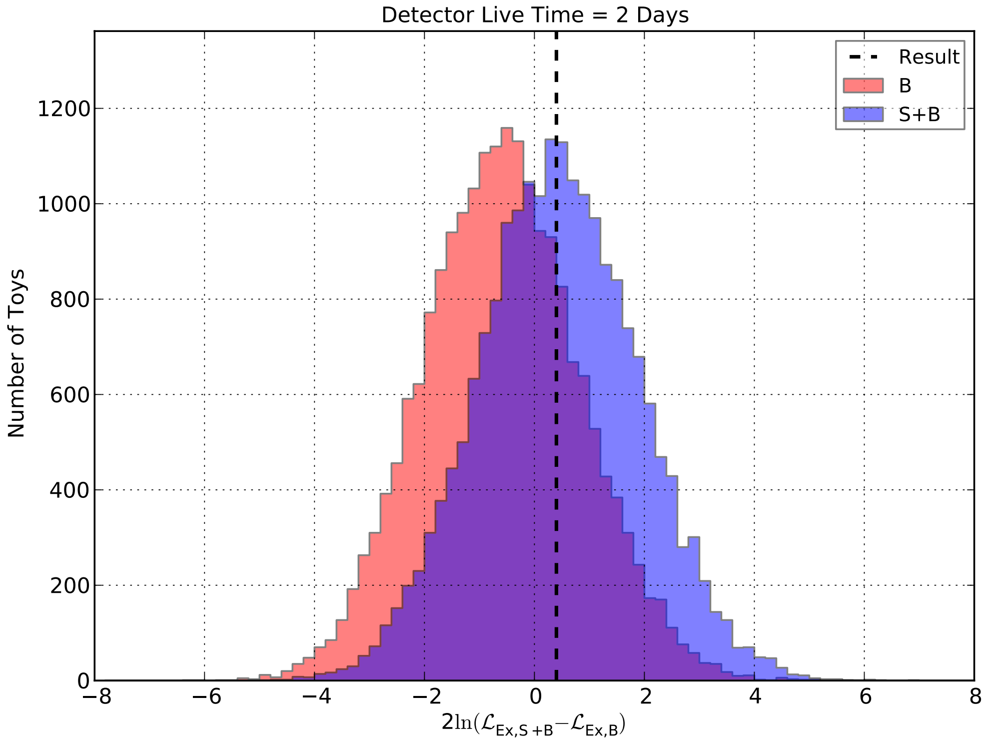

Figure 6.35: Simulated distributions of the extended likelihood ratio test statistic - no significant preference for either hypothesis is observed.

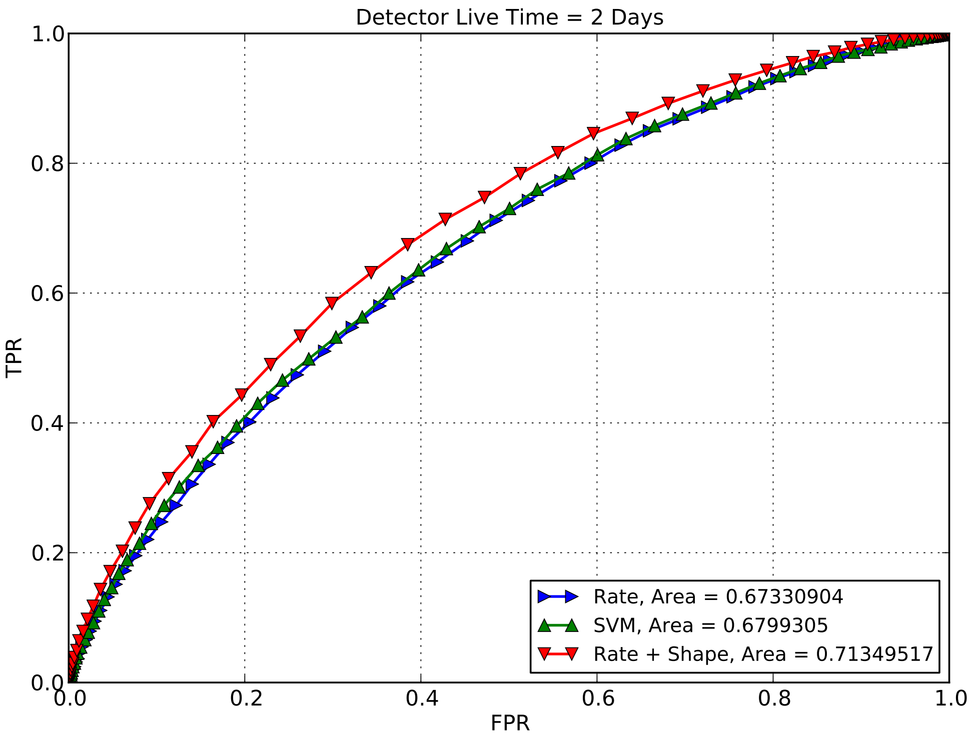

Figure 6.36: ROC curves for the three different test statistics used to search for the neutrino signal. The extended likelihood ratio test (using both shape and rate information) provides the best sensitivity.

of the three tests, suggesting this test is the most sensitive. The next most sensitive test is the rate test that uses the SVM classifier, which performs slightly better than the rate test that uses the cut-based selection criteria.

6.7 Conclusions

A full IBD analysis has been performed to search for a neutrino signal at the large scale SoLid prototype module for the first time. Several techniques have been explored to define a signal-to-background rich region of the detector phase space, using traditional cut-based methods, as well as machine learning techniques, including support vector machines and boosted decision trees. The cut-based techniques provide a background reduction of around 200 for an acceptable analysis efficiency of 57%, and the machine learning techniques explored (once tuned) are able to reduce the background by an additional 20% for the equivalent signal efficiency. The topological variables provide the largest overall background reduction of those explored.

Several statistical tests have been performed to compare the reactor on and off datasets, using both rate and shape information, and using cut-based selection criteria as well as machine learning. No statistically significant neutrino signal has been observed due to the limited size of the dataset - the data is compatible with both a background only and a signal plus background hypothesis. The results show that the background model constructed for the reactor on period, using data driven techniques, is appropriate. Studies using simulated signal events indicate that the extended likelihood ratio test, which uses both rate and shape information, is the most sensitive.

These techniques and methods will be re-used once the first data from the first phase of SoLid is available. This work has also informed the design and optimisation of the full experiment, to overcome the limitations of this analysis. A selection of these improvements are described in chapter 7. As well as background reduction, another key parameter for performing a sterile neutrino oscillation search is energy resolution. The next chapter shows how cosmic muons can be used to calibrate the energy response of the detector.

Chapter 7

Phase I Outlook

7.1 Introduction

This short chapter describes a selection of optimisations of the phase 1 detector design, informed by experience from the prototype, including many results from the previous chapters. Two elements of the design are discussed: passive shielding and triggering. These two topics have been chosen since they aim to resolve the limitations found in the IBD analysis presented in chapter 6. Section 7.2 describes how the IBD correlated background can be tackled further by the use of passive water shielding. It was shown in chapter 3 that the choice of trigger scheme for the prototype resulted in a significantly reduced neutron trigger efficiency - section 7.3 describes on-going studies of new trigger algorithms that are dedicated to neutron signals. Finally, given the gains in signal efficiency and background reduction from these upgrades, the expected physics reach of phase 1 of the experiment is presented in section 7.4.

Many of the results shown in this section are the result of collaborative efforts between the author and members of the SoLid collaboration - where appropriate, this is specified.

7.2 Passive Shielding

It was shown in chapter 6 that the dominant IBD background contribution, post offline selection criteria, is the correlated background. Sources of this background include fast cosmic neutrons in the energy range of 4-15 MeV. Above 15 MeV, the neutrons are not expected to thermalise whilst crossing the detector [71]. Fast neutrons may also be produced as decay products of cosmic muon interactions (i.e. muon spallation). Simulations of these cosmic backgrounds have been used to study the effect of passive shielding for phase 1 (recall that the prototype had no passive sheidling). Two cosmic radiation generators are used: **Guan**, which simulates cosmic muons; and **Gordon**, which simulates cosmic neutrons. Comparisons of the output of these generators with background data from the large scale prototype shows reasonable agreement [71].

Materials with high hydrogen density, such as water, paraffin or HDPE plastic, are suitable neutron moderators. Neutrons are attenuated due to elastic scattering by the hydrogen protons [72]. Taking into account costing, it has been decided that the entire detector will be surrounded by a 50 cm water shield wall - see figure 7.1. The top of the detector will be covered by 50 cm of HDPE plastic. The simulated energy distribution of neutrons, before and after the use of this shielding, is shown in figure 7.2. It can be seen that in the energy range of fast neutrons (4-15 MeV), the flux reduces by around an order of magnitude.

Integrating over this region, the rate of fast neutrons is expected to decrease by $\sim 90\%$, and the total reduction in neutron rate is $\sim 60\%$ [71]. Since comparisons between simulation and data have shown that these neutrons are the dominant source of correlated background, this background contribution is expected to be reduced by a similar fraction compared to the prototype. The attenuated neutrons may increase the relative rate of thermal neutrons (< 250 keV), which can still contribute to the accidental background, although as shown in chapter 6, this is not as dominant as the correlated contribution.

7.3 Triggering

Recall that the large scale prototype uses a single trigger configuration - a threshold-based trigger, requiring at least 2 channels to be above threshold within a time-window of $\pm$ 3 samples. It was shown in chapter 3 (figure 3.9) that the trigger threshold had to be placed very high in order to achieve a manageable data rate for the 2015 data-taking runs, resulting in a neutron trigger efficiency of only $\sim 5\%$. Further, the prompt and delayed events forming IBD candidates are required to trigger separately, meaning low amplitude signals (< 500 keV) in co-incidence with these events are not recorded. This could include, for example, proton recoil events, which can in principle be used as a tag of the fast neutron correlated background.

This is rectified for phase 1 by using a dedicated neutron trigger. This trigger is combined

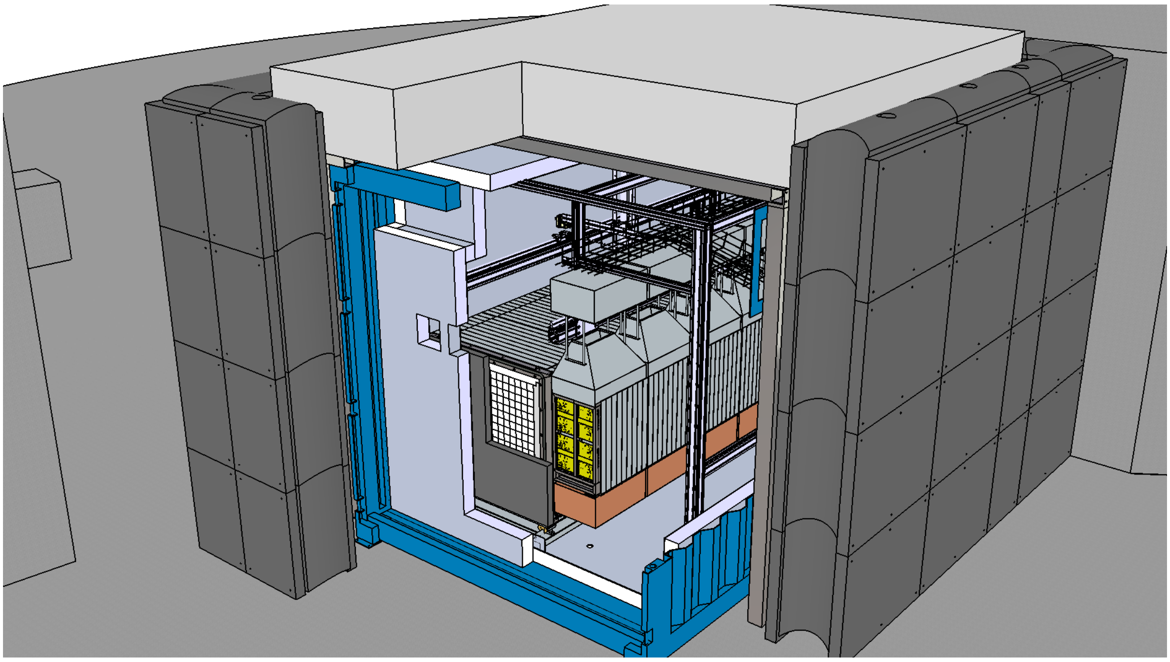

Figure 7.1: Drawing of phase 1 of SoLid, including a 50 cm water shield wall (composed of bricks), and HDPE roof.

with a buffer, such that a neutron trigger will also read-out the prompt event data. Multiple neutron identification algorithms that can be implemented in the phase 1 electronics have been tested in software, using data from a lab setup. The setup uses a single cube with two SiPMs, and additionally, a photo-multiplier tube (PMT) facing one of the cube sides with a hole in the Tyvek wrapping - see figure 7.3. The PMT is able to detect significantly more light than that extracted by the fibres, and can perform highly pure and efficient neutron identification [54]. The following provides a description of the neutron ID algorithms explored:

- **Time over Threshold (ToT)**: the number of samples above a pre-set threshold is counted within a rolling time-window. Since the scintillation time of neutron signals from the ZnS is slower than that of EM signals in the PVT, the ToT is larger for neutrons.

- **Maximum Amplitude**: a threshold requirement is set for the amplitude of the waveforms. This is similar to the requirement used for the large scale prototype (without the co-incidence requirement), and is not especially suitable for neutron identification.

- **Integral/Amplitude**: this is the algorithm used to identify neutrons during offline analysis of the prototype data, as described in chapter 4 (without the co-incidence requirement). This algorithm is too complex to be implemented in the phase 1

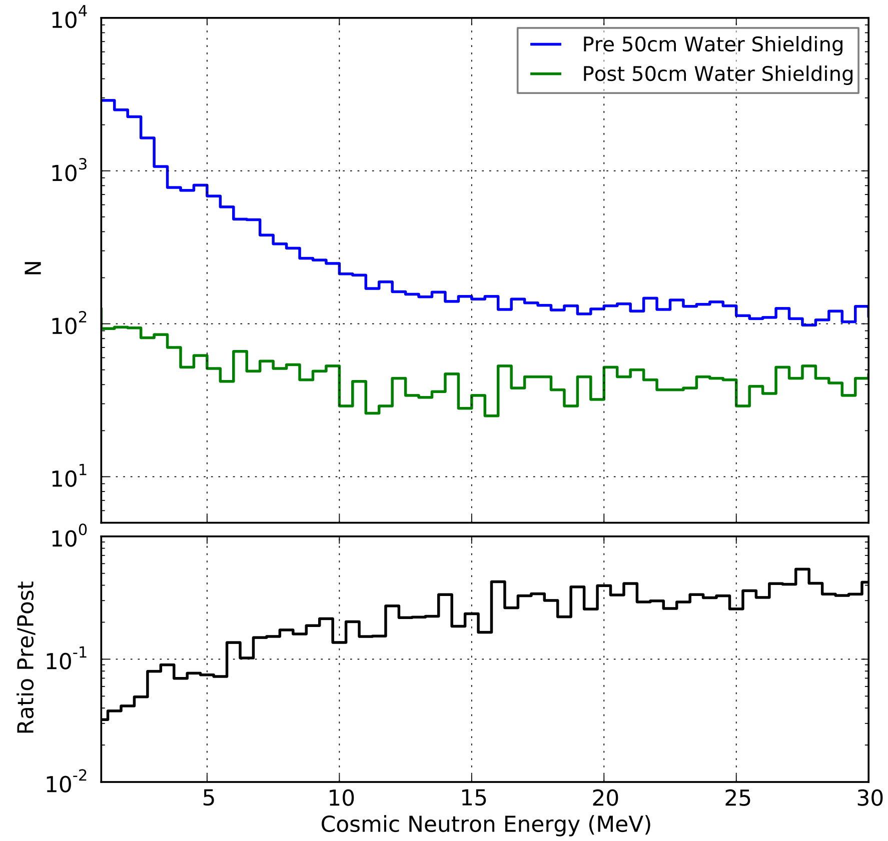

Figure 7.2: Cosmic neutron energy distribution using the Gordon cosmic generator, with and without the use of passive water shielding.

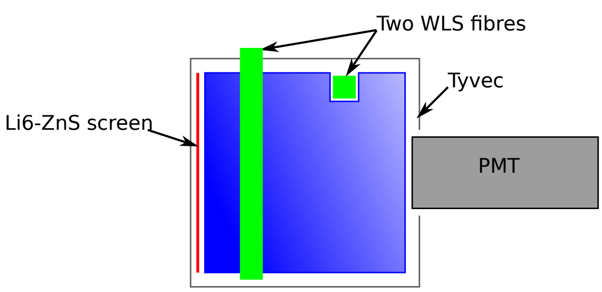

Figure 7.3: Diagram of a lab setup which uses a single cube, two silicon photo-multipliers and a photo-multiplier tube (PMT). The PMT is able to perform very effective neutron identification due to the direct light extraction (i.e. no optical fibres).

electronics, and is included only for reference.

- **Number of Pixel Avalanches**: the integral of the waveform in units of PA.
- **Number of Peaks**: the number of peaks (i.e. local maxima) within a time window. As seen in neutron waveform examples, the ZnS signals can give rise to multiple sub-peaks in the waveforms.

The ROC curves showing the efficiency vs. false positive rate of these different algorithms are shown in figure 7.4. To make this plot, waveforms were input into the algorithms on a rolling basis, to emulate the electronics trigger. It can be seen that many of the phase 1 candidate algorithms, such as 'Number of Peaks', are far more efficient than that used for the prototype (for the equivalent false positive rate). The performance of these upgrade algorithms is similar to the prototype offline reconstruction algorithm. The distributions of the 'Maximum Amplitude' and 'Number of Peaks' parameters are shown in figure 7.5. Far greater separation is observed for the phase 1 algorithm than that used for the prototype. The algorithms 'Time over Threshold' and 'Number of Peaks' are currently being implemented in the FPGA firmware of the phase 1 electronics. The target neutron trigger efficiency for phase 1 is $\sim 80\%$.

7.3.1 Prompt Detection

Upon triggering, a zero-suppressed buffer that stores data from certain channels for the previous $\sim 500\,\mu$s is read-out. This is designed to contain data from the prompt event [54]. The channels read-out can be from ± 2 planes either side of the plane that registered the neutron. This removes the requirement for the prompt event to trigger separately, and also allows any small signals that are co-incident with the prompt event to be detected with high efficiency. Other trigger configurations will also be used simultaneously, including a random trigger and threshold-based trigger for EM signals above $\sim 1\,$MeV (e.g. to read-out muons).

The trigger efficiency for the prototype, and that expected for phase 1 from simulation, is shown as a function of visible EM energy in figure 7.6. The plot was made by simulating EM events in a single cube, over a range of energies. Specifically:

- A range of EM energies were scanned:

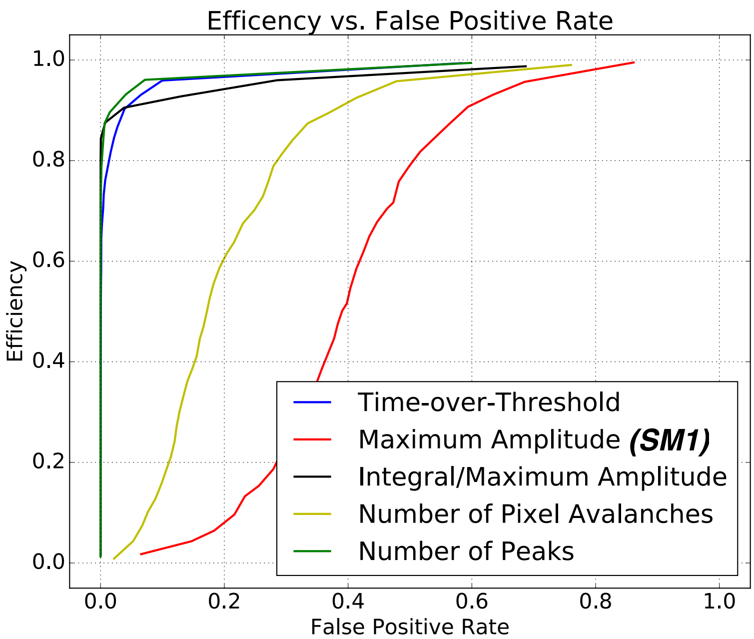

Figure 7.4: ROC curves for potential phase 1 neutron trigger algorithms, using a single cube lab setup. A single SiPM channel is used for the analysis. The large scale prototype configuration is equivalent to the maximum amplitude curve. An AmBe source was used to provide neutrons, and so the false positive rate will be different for SoLid. *Credit: Lukas Arnold.*

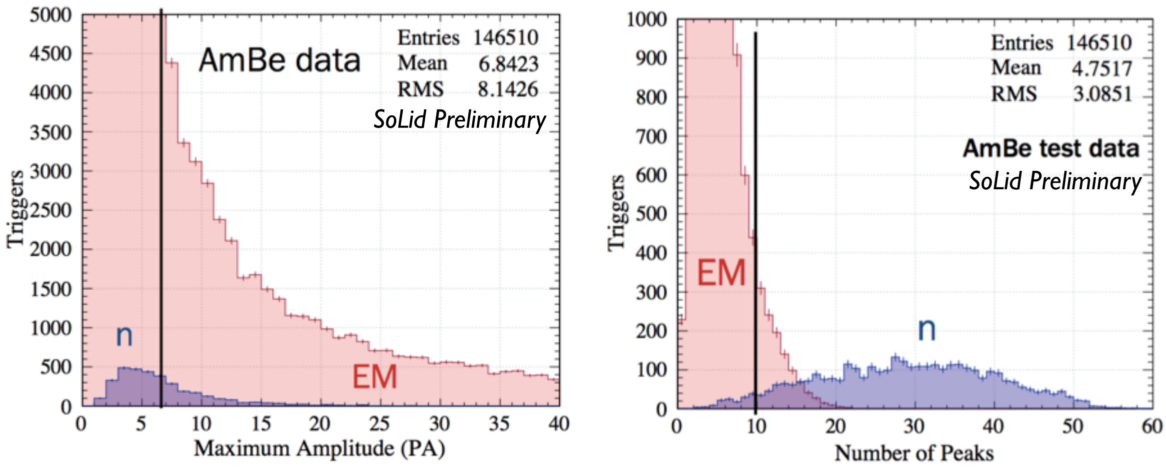

Figure 7.5: Comparison of two neutron identification parameters (described in the text) that can be implemented in the phase 1 electronics. Left: the maximum amplitude, where the black line represents the threshold used for the large scale prototype. Right: Number of Peaks, where the black line shows a proposed threshold. The corresponding ROC curves are shown in figure 7.4. *Credit: Lukas Arnold.*

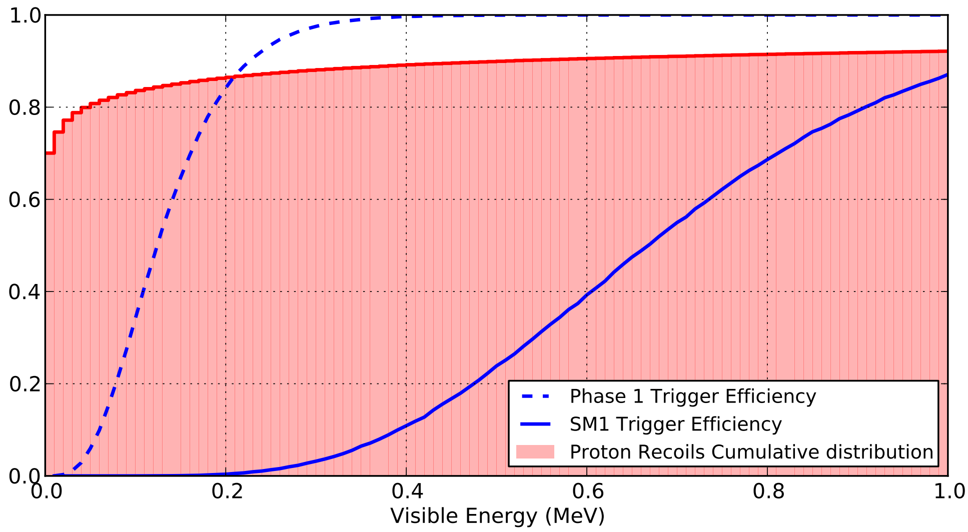

Figure 7.6: Trigger efficiency as a function of visible EM energy using simulation. The phase 1 trigger configuration is described in the text. The visible energy distribution of proton recoils is found using the Gordon cosmic neutron generator.

- At each step, the light yield corresponding the to the energy is calculated (recall that the light yield of phase 1 is increased by $\sim 20\%$ per fibre compared to the prototype).
- Generate 10k EM events for the cube:
  * Simulate the number of photons transmitted down the fibres. To do this, a Poisson distribution is sampled for each optical fibre coupled to the cube, whose mean is the light yield. Note: phase 1 has twice as many optical fibres.
  * Apply the trigger conditions of both detector setups, to simulate whether the event would be read-out (via the neutron trigger buffer).

Figure 7.6 also shows the cumulative energy distribution of proton recoils from cosmic neutrons, and it can be seen that the efficiency of detecting these proton recoils should increase by around 5%. Chapter 6 described how fast cosmic neutrons can be tackled via selection criteria based on the multiplicity of the prompt event. The lowered threshold of the visible energy of phase 1 implies the rejection power of this selection criteria will increase. The multiplicity of protons for the two setups is shown in figure 7.7 using the Gordon generator for cosmic neutrons. Here, the multiplicity is the number of protons that are time co-incidence (within 100ns) of each-other. It was also required that at least one of the protons have a visible energy above 1 MeV (the target prompt energy selection cut for phase 1). It can be seen that the phase 1 distribution is spread towards higher values. This may allow fast neutron events to be tagged with improved efficiency for phase 1.

7.4 Sterile Search Sensitivity

The sensitivity of the phase 1 detector to discover sterile neutrinos has been studied using simulation. Chapter 2 outlined how oscillations at the very short baseline would be driven by the Δm^2_{14} and θ_{ee} parameters. Since the oscillations depend on E/L, oscillations distort the energy spectrum of neutrinos as a function of distance from the reactor. An extreme example is shown in figure 7.8 - the no-oscillation hypothesis is vastly different to the sterile oscillation hypothesis. A 2D E vs. L binned marginalised χ^2 comparison is performed to compare the experimental data with the no oscillation hypothesis - see [73] for more details. Only shape information is used in the comparison, not absolute rate, to remove dependency from both the theoretical predicted rate and the reactor power (this is comparable to current

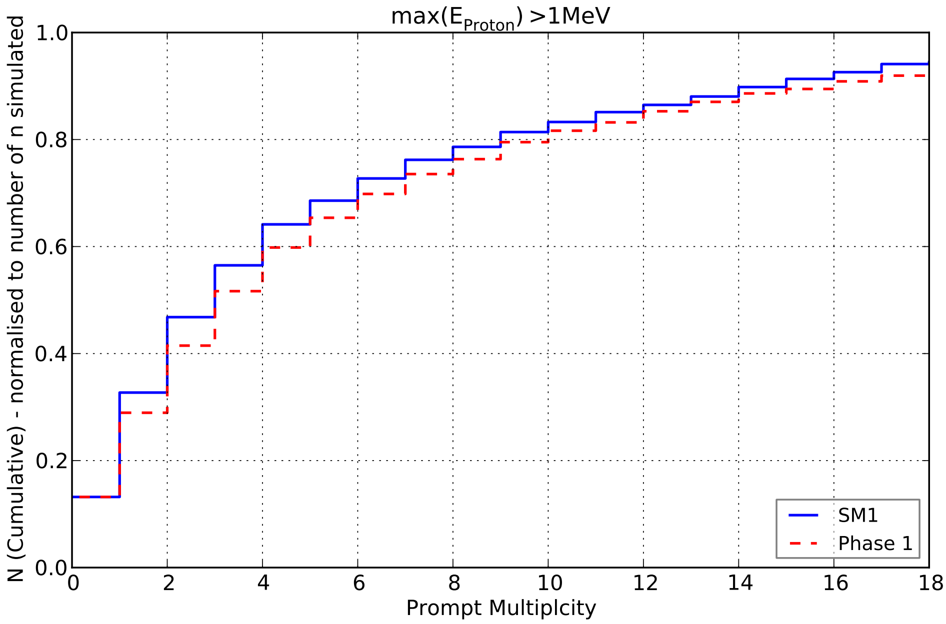

Figure 7.7: Multiplicity of prompt events from simulated background IBD candidates (using the Gordon cosmic neutron generator) N.b: not all selection criteria outlined in chapter 6 has been applied due to low statistics - those that have include: trigger, Δt, $E_{Prompt} > 1$ MeV. *Made in collaboration with Yamiel Abreu.*

experiments using a near and far detector at larger baselines).

A scan over the two driving oscillation parameters has been performed, and the regions where the detector is sensitive to the discovery of a sterile neutrino (3+1 model) is shown in figure 7.9. The allowed regions (95 % CL) from the reactor and gallium anomalies are also shown, as is the best fit from Kopp [2] (using short baseline experiments only). The exclusion contour shows the boundary where the value of the χ^2 drops below the 95 % confidence limit. For example, in regions on the lower-left of the plot, the sterile neutrino signal can not be distinguished from the non-sterile hypothesis with high certainty; whereas a sterile neutrino on the mid-right region would distort the E vs L distribution so much that a χ^2 comparison with the non-sterile hypothesis would return a very low p value (suggesting some alternative hypothesis is more appropriate).

It can be seen that the best fit point from Kopp et. al. can be excluded to 95% confidence level after 150 days of reactor on running (corresponding to around 1 year of operation at the BR2 reactor). The input parameters to this search are shown in the figure, where the signal

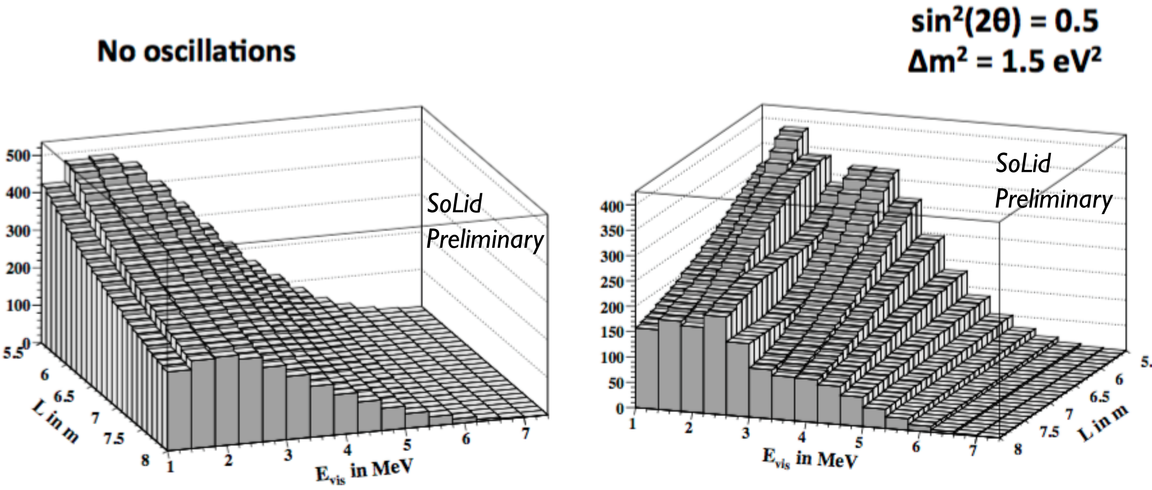

Figure 7.8: Example distribution of E vs L for two oscillation hypothesis. Left: no sterile oscillation case. Right: an extreme 3+1 oscillation.

to background ratio S/B (AKA S/N), absolute IBD detection efficiency ϵ_{IBD} and mass are the targets of the experiment.

7.5 Conclusions

Phase 1 of the SoLid experiment will search for sterile neutrinos in the region of phase space around the reactor anomaly. The current best fit values for a 3+1 sterile neutrino, found from the analysis of results from previous short baseline neutrino experiments (Kopp et. al.), should be confirmed or refuted after around 1 year of data taking to 95% CL. A marginalised χ^2 test is used to compare the sterile and non-sterile hypothesis. To achieve it's aims, the experiment is required to have a signal to background ratio of around 3:1, with an absolute neutrino signal efficiency of around 30% (post analysis). To reach these targets, the phase 1 detector has been upgraded compared to the large scale prototype, many of which are motived by work presented in the previous chapters. A dedicated neutron trigger, using more sophisticated algorithms at the FPGA level, similar to those used for the prototype offline reconstruction, will increase the trigger efficiency from $\sim 5\%$ to a target of $\sim 80\%$. A 50 cm water wall and 50 cm HDPE roof will be used for to moderate cosmic fast neutrons. Simulation results suggest the correlated background in the IBD analysis can be reduced by around an order of magnitude.

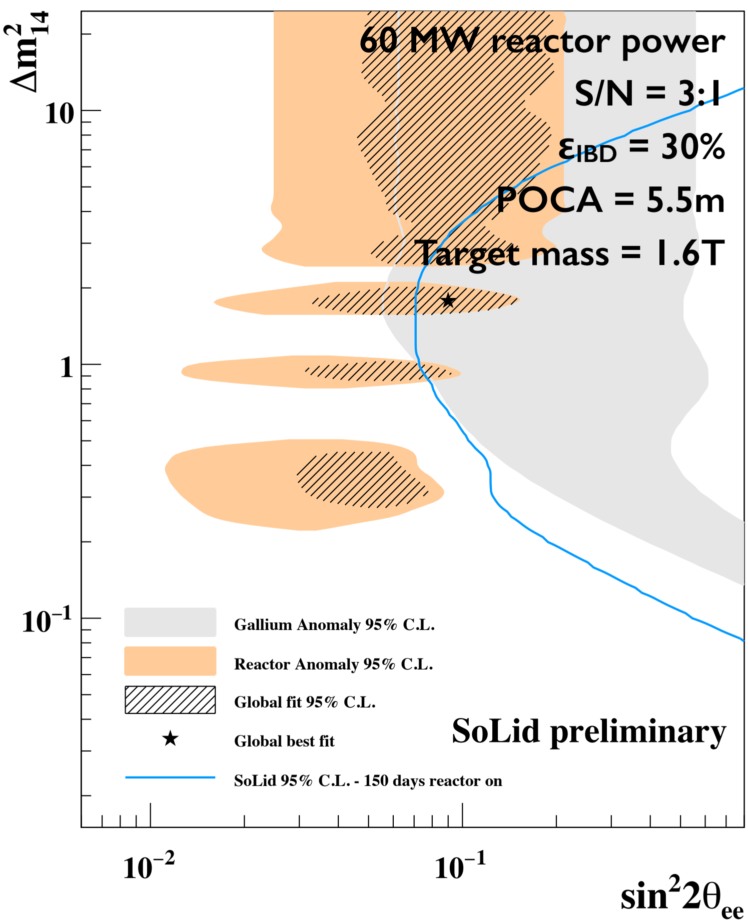

Figure 7.9: Sensitivity of SoLid phase 1. The best fit point for the 3+1 sterile hypothesis can be excluded after 150 days of reactor on running. *Credit: Leonidas Kalousis [73].*

Chapter 8

Conclusion

The SoLid experiment is one of the next generation of very short baseline reactor neutrino experiments. These experiments will search for sterile neutrinos at baselines $\mathcal{O}(10)$m, using compact research reactors as a source of anti-electron neutrinos. Neutrino oscillations at this distance scale could help resolve many anomalies observed by previous experiments, including the reactor and gallium anomalies, which are both significant to around the 3σ level. By studying different core fuel compositions, they can also inform studies of the nuclear fission cross sections, such as resolving the source of the 5 MeV distortion reported by the current generation of short baseline reactor experiments. Historically, it was anomalies similar to these, such as the solar neutrino problem, that led to the discovery of neutrino oscillations; and hence, the conclusion that neutrinos are infact massive. These future experiments will be placed closer to reactors than ever before. This introduces new challenges, in particular extreme environmental backgrounds. SoLid uses a novel technology to combat these challenges. The entire detector is formed of $5\,\mathrm{cm}^3$ cubes that can detect both neutrons and electromagnetic signals. The viability of the technology has been demonstrated by a large scale 288 kg prototype that was constructed and commissioned in early 2015.

The first event reconstruction of data recorded by the large scale prototype has been performed. Many reconstruction algorithms have been developed for event finding, event filtering, tracking, and calibration of the raw detector data. These algorithms, and their software implementation, form the blueprint of how reconstruction will be implemented for phase 1 of the experiment. This includes methods to identify and accurately reconstruct positrons and neutrons, which together form the neutrino signal; and muons, which are a background.

A full inverse beta decay analysis has been performed. The neutrino signal distributions have been studied using simulation, including readout and reconstruction effects, and the backgrounds have been studied using data, taking advantage of the reactor off signal-free datasets. Different techniques, including traditional cut-based as well as machine learning methods, have been applied to find a first selection of neutrino IBD candidates. Using these techniques, the extreme backgrounds of the neutrino signal were reduced by around a factor of 200. The accidental background, which is the largest component, can be reduced by a factor around 1000. The most effective variables for distinguishing signal and background candidates are the topological features, such as the radial separation between the neutron and positron candidates, and the positron multiplicity. Machine learning techniques, including support vector machines (SVMs) and boosted decision trees (BDTs), are both able to further reduce the background by $\sim 20\%$. A background model for the reactor on period has been found using data driven techniques. Comparisons of the reactor on data with this model have shown the reactor on dataset is statistically limited to confirm the presence of a neutrino signal, whilst demonstrating the background has been modelled accurately. These comparisons used a variety of tests, using both rate and shape information, such as the prompt energy distribution.

The energy response of the detector has been calibrated using cosmic muons. Muons can be effectively separated from other kinds of events based on energy deposited and topological information. By selecting muons that trigger many channels simultaneously, the path of the muons can be accurately reconstructed, allowing the energy deposited per unit path length dE/dx to be reconstructed. This parameter is a standard candle. Geometrical selection criteria have been developed to select cases where the energy deposited can be localised to individual cubes, allowing the energy response of individual cubes to be equalised to a precision of 2%. From comparisons with simulation, the scale of the energy response of the detector can be found to a similar precision.

This work has informed many design improvements for the phase 1 detector compared to the prototype. A dedicated trigger will be used to perform neutron identification at the electronics trigger level. This will significantly increase the detection efficiency of neutrons, whilst maintaining a manageable data-rate. By combining this trigger with a time buffer, the prompt event can also be read-out, removing the requirement of prompt events to cause a trigger. This allows for much more data from the prompt event to be read-out compared to the prototype. The correlated background, which is the dominant source of background after IBD selection criteria have been applied, will be significantly reduced with passive shielding. A 50 cm water wall and HDPE plastic roof will be constructed. The first oscillation search will be performed within one year of data taking by the 1.6 tonne phase 1 detector, by which time SoLid aims to be able to accept or reject the best fit $3 + 1$ sterile neutrino hypothesis to 95% confidence level.

Abbreviations List

BDT: Boosted Decision Tree.

BR2: Belgian Research Reactor 2.

CPT: Charge, Parity, and Time Reversal Symmetry.

HFRI: High Flux Isotope Reactor.

IBD: Inverse Beta Decay.

ILL: Institut Laue-Langevin.

ML: Machine Learning.

MPV: (Landau) Most Probable Value.

PA: (SiPM) Pixel Avalanche.

PDE: Photon Detection Efficiency.

PMNS: Maki, Nakagawa, Sakata and Pontecorvo (marix).

ORNL: Oak Ridge National Laboratory.

SiPM: Silicon Photomultiplier.

SM: Standard Model.

SoLid: Short baseline Oscillation search using a Lithium-6 Detector (experiment).

SVM: Support Vector Machine.

Bibliography

[1] Patrignani et al (Particle Data Group). (2016). Phys C **40** 100001.

[2] Kopp, J. (2013). *Sterile Neutrino Oscillations: The Global Picture*. **arXiv:1303.3011**.

[3] TODO. *TODO*. **TODO**.

[4] Vacheret, A (SoLid Collaboration). (2016). *A novel solid segmented scintillator antineutrino detector*. JINST.

[5] Griffiths, D. (2008). *Introduction to Elementary Particles*. Wiley.

[6] Waller, D (SNO Collaboration) (2004). *Results from the Sudbury Neutrino Observatory*. SLAC Summer Institute on Particle Physics, SSI04.

[7] Raaj, J (Super-K Collaboration). (2008). *Solar and Atmospheric Nuetrinos in Super-Kamiokande*. Neutrino 2008 Conference.

[8] Davis et al. (1968). *Search for Neutrinos from the Sun*, Phys Rev Lett **20**, 1205.

[9] Drewes, M. (2013). *The Phenomenology of Right Handed Neutrinos*. **arXiv:1303.6912**.

[10] Hirata et al (KAMIOKANDE-II). (1988). Phys Lett **B205**, 416.

[11] Giunti and Kim. (2007). *Neutrino Physics and Astrophysics*. Oxford University Press.

[12] Dell'Oro et al. (2016). *Neutrinoless double beta decay: 2015 review*. **arXiv:1601.07512**.

[13] Fukuda et al (Super-Kamiokande Collaboration). (1998). *Evidence for Oscillation of Atmospheric Neutrinos*. Phys Rev Lett **81**, 1562.

[14] Pontecorvo, B. (1968). *Neutrino Experiments and the Problem of Conservation of Leptonic Charge*. Sov Phys JETP **26**, 984 - 988.

[15] Bian, J. (2013). *The NOvA Experiment: Overview and Status*. **arXiv:1309.7898**.

[16] Barabash, A (SuperNEMO Collaboration). (2011). *SuperNEMO double beta decay experiment*. **arXiv:1112.1784**.

[17] T2K Collaboration. (2011). *The T2K Experiment*. **arXiv:1106.1238**.

[18] KamLAND Collaboration. (2002). *First Results from KamLAND: Evidence for Reactor Anti-Neutrino Disappearance*. **arXiv:hep-ex/0212021**.

[19] Cao et al. (2016). *An overview of the Daya Bay Reactor Neutrino Experiment*. **arXiv:1605.01502**.

[20] Daya Bay Collaboration. (2015). *A new measurement of antineutrino oscillation with the full detector configuration at Daya Bay*. **arXiv:1505.03456**.

[21] Ahn et al (RENO Collaboration). (2010). *RENO: An Experiment for Neutrino Oscillation Parameter θ_{13} Using Reactor Neutrinos at Yonggwang*. **arXiv:1003.1391**.

[22] Suekane et al (Double Chooz Collaboration). (2016). *Double Chooz and a History of Reactor Theta13 Experiments*. **arXiv:1601.08041**.

[23] Borexino Collaboration. (2013). *SOX: Short distance neutrino Oscillations with BoreXino*. **arXiv:1304.7721**.

[24] Bilenky, S. (2010). *Introduction to the Physics of Massive and Mixed Neutrinos*. Springer.

[25] Parke, S. (2016). *What is Δm^2_{ee}?*. **arXiv:1601.07464**.

[26] Gando et al. (KamLAND Collaboration). (2013). *Reactor On-Off Antineutrino Measurement with KamLAND*. Phys Rev D **88**, 033001.

[27] T2K Collaboration. (2011). *Measurements of the T2K neutrino beam properties using the INGRID on-axis near detector*. **arXiv:1111.3119**.

[28] Tankaka (T2K Collaboration). (2016). *T2K: Latest Results*. Neutrino 2016 Conference.

[29] Petcov, S. (2014). *Leptonic CP Violation and Leptogenesis*. **arXiv:1405.2263**.

[30] Matthew Murdoch, *A Safeguards Friendly Device for Monitoring Reactor Anti-Neutrinos with T2K Technology*, `http://indico.hep.manchester.ac.uk/getFile.py/access?contribId=203&sessionId=6&resId=0&materialId=slides&confId=4534`

[31] Saint-Gobain. (2016). *Scintillating Optical Fibres*. Saint-Gobain Catalouge

[32] Reines and Cowan. (1956). *Detection of the Free Neutrino: A Confirmation*, Science Vol 124, 3212, 103-104.

[33] Mueller et al. (2011). *Improved Predictions of Reactor Antineutrino Spectra*. **arXiv:1101.2663**.

[34] Mention et al. (2011). *The Reactor Antineutrino Anomaly*. **arXiv:1101.2755**.

[35] Joo, K. (2016). *Results from RENO and prospects with RENO-50*. Neutrino 2016 Conference.

[36] Huber, P. (2016). *The 5 MeV bump - a nuclear whodunit mystery*. **arXiv:1609.03910**.

[37] Giunti and Laveder. (2010). *Statistical Significance of the Gallium Anomaly*. **arXiv:1006.3244**.

[38] Bowden, N. (2016). *Review of short-baseline reactor antineutrino experiments*. Neutrino 2016 Conference.

[39] Helaine, V. (2016). *Sterile neutrino search at the ILL nuclear reactor: the STEREO experiment*. **arXiv:1604.08877**.

[40] Hamamatsu SiPM datasheet: *http://www.hamamatsu.com/eu/en/S12572-050P.html*

[41] Donzella, A. (2010). *A proton recoil telescope for neutron spectroscopy*. Nucl Instrum Meth Phys Res A **613**, 58-64.

[42] Morrison, I. (2008). *Introduction to Astronomy and Cosmology*. Wiley.

[43] The Kamland Collaboration. (2010). *Production of radioactive isotopes through cosmic muon spallation in KamLAND*. **arXiv:0907.0066**.

[44] Matplotlib plotting package: *http://matplotlib.org*

[45] Simon Vercaemer. *Neutron identification in the SoLid experiment*. July 2015, EPS Proceedings.

[46] Lasserre, T. *ECT Lecture 2 - Reactor Antineutrino Detection*.

[47] Cowan, G. *Computing and Statistical Data Analysis*. University of London Postgraduate Lectures.

[48] Muon Spallataion thing from Fred. *Lorem Epsum...*. Location and date.

[49] Kalousis et al (SoLid Collaboration). (2016). *Neutron detection efficiency in SM1 (internal tech. note)*. SoLid doc-db: 371-v7.

[50] Munn, P (2016). *Background mitigation in the PROSPECT short-baseline reactor experiment and other surface detectors*. Neutrino 2016 Conference.

[51] Saunders, D (2013). *Visualising 5D data at LHCb*. Bristol University (MSci Thesis).

[52] Scikit learn classifier comparison: *scikit-learn.org*

[53] Ryder, N. (2015). *First results of the deployment of a SoLid detector module at the SCK-CEN BR2 reactor*. **arXiv:1510.07835**

[54] Ryder, N (SoLid Collaboration). (2016). *SoLid anti-Neutrino Detector Read-Out System*. IEEE 2016 Conference.

[55] Kalousis, L (SoLid Collaboration). (2016). *SoLid : A compact detector for very shortÂŋâĂŘbaseline neutrino experiments*. ICHEP 2016 Conference.

[56] Jianming Bian (Noνa Collaboration). (2015). *First Results of ν_e Appearance Analysis and Electron Neutrino Identification at NOvA*. **arXiv:1510.05708**

[57] Oak Ridge National Laboratory. (2002). *Monte Carlo N-Particle Transport Code System for Multiparticle and High Energy Applications* (MCNPX Manual).

[58] Barlow, R. (1999). *Statistics: A Guide to the Use of Statistical Methods in the Physical Sciences*. Wiley. p40.

[59] Keck, T. (2016). *Multivariate Classification and Machine Learning in HEP*. Inverted CERN School of Computing 2016.

[60] Neyman and Pearson. (1933). *On the Problem of the Most Efficient Tests of Statistical Hypotheses*. **DOI: 10.1098/rsta.1933.0009**.

[61] Cortes and Vapnik. (1995). *Support-Vector Networks*. **DOI: 10.1007/BF00994018**.

[62] Freund and Schapire. (1997). *A decision-theoretic generalization of on-line learning and an application to boosting*. Journal of Computer and System Sciences **55(1)**, 119âĂŞ139.

[63] Boursette, D (SoLid Collaboration). (2016). *Light yield and energy resolution studies for SoLid phase 1*. Neutrino 2016 Conference.

[64] Agostinelli et al (Geant4). (2003). *A simulation toolkit*. Nucl Instrum Meth Phys Res A **506(3)**, 250-303.

[65] Vogel and Beacom. (1999). *Angular distribution of neutron inverse beta decay*. Phys Rev D60 053003.

[66] Barlow, R. (1990). *Extended maximum likelihood*. Nucl Instrum Meth **A297**, 496-506. **DOI: 10.1016/0168-9002(90)91334-8**.

[67] Muiden, J. (2016). *Determination of the Accuracy and Precision of BR2*. SoLid doc-db: 256-v1.

[68] Ghys, L. (2016). *Fine-tuning of MPPC parameters in readout simulation (collaboration meeting presentation)*. SoLid doc-db: 316-v1.

[69] Guan, M. (2015). *A parametrization of the cosmic-ray muon flux at sea-level*. **arXiv:1509.06176**.

[70] Barrand G et al. (2000) *GAUDI - A software architecture and framework for building LHCb data processing applications*. CHEP 2000 Conference.

[71] Pinera, I (SoLid Collaboration). (2016). *Identification of Background Components with the SoLid anti-Neutrino Detector*. IEEE 2016 Conference.

[72] Dobrzynski, L.; K. Blinowski (1994). *Neutrons and Solid State Physics*. Ellis Horwood Limited. ISBN 0-13-617192-3.

[73] Kalousis, L (SoLid Collaboration). (2016). *The SoLid experiment*. Neutrino 2016 Conference.