A search for *hep* solar neutrinos at the Sudbury Neutrino Observatory

Timothy J. Winchester

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2016

Reading Committee:

Nikolai Tolich, Chair

Jason Detwiler

Miguel Morales

Program Authorized to Offer Degree:
Physics

University of Washington

Abstract

A search for *hep* solar neutrinos at the
Sudbury Neutrino Observatory

Timothy J. Winchester

Chair of the Supervisory Committee:
Affiliate Professor Nikolai Tolich
Physics

Solar neutrinos from the fusion *hep* reaction, $^{3}\mathrm{He} + p \rightarrow {}^{4}\mathrm{He} + e^{+} + \nu$, have previously remained undetected due to their flux being about one one-thousandth that of ^{8}B neutrinos. These neutrinos are interesting theoretically because they are less dependent on solar composition than other solar neutrinos, and therefore provide a somewhat independent test of the Standard Solar Model. In this analysis, we develop a new event fitter for existing data from the Sudbury Neutrino Observatory. We also use the fitter to remove backgrounds that previously limited the fiducial volume, which we increase by 30%. We use a modified Wald-Wolfowitz test to increase the amount of live time by 200 days (18%) and show that this data is consistent with the previously-used data. Finally, we develop a Bayesian analysis technique to make full use of the posterior distributions of energy returned by the event fitter. In the first significant detection of *hep* neutrinos, we find that the most-probable rate of *hep* events is $3.5 \times 10^{4}/\mathrm{cm}^{2}/\mathrm{s}$, which is significantly higher than the theoretical prediction. We find that the 95% credible region extends from 1.0 to 7.2 $\times 10^{4}/\mathrm{cm}^{2}/\mathrm{s}$, and that we can therefore exclude a rate of 0 *hep* events at greater than 95% probability.

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

Table Number Page

ACKNOWLEDGMENTS

This dissertation is built on the work of many people and would not have been possible without the help and understanding of many more.

First, I'd like to thank the SNO collaboration for all the effort they put into building a cutting-edge experiment, collecting years of data, and curating the data into something ready for this analysis. It was kind of you to let me join the club.

I'd also like to thank the folks at CENPA for being friendly, knowledgeable, and willing to put in time to help with technical issues of all sorts. In particular, Gary Holman was always ready to tackle issues relating to computing resources.

I've made use of many tools that were made publicly available by the U.S. and international science communities. These include ROOT, Matplotlib [24], the Centre for Advanced Computing (formerly HPCVL), and the Rocks (formerly Athena) cluster at CENPA. Without these powerful tools at my disposal, this work would have suffered greatly. I'd also like to mention the people who have worked to make C++, Python, NumPy, and SciPy the powerful and free tools that they are. Thank you.

Several undergraduates have contributed to my work in grad school, both the scintillator scattering measurements and the *hep* analysis. Specifically, Marija Glisic, Alex Huss, Andrew Eberhardt, Richard Hong, and Richard Seng all spent time behind a computer screen or in the lab (or both) to make my life easier.

Many other people at the University of Washington have also been important to my success in graduate school. Two I'd like to mention specifically are Catherine Provost and Nikolai Tolich. Catherine has been a helpful and dependable confidante and adviser. Nikolai has done much more than being a thesis adviser: he's been an advocate, mentor, and role

model.

Finally, I would like to thank Gary. You've been so supportive all these years and I don't know what I'd do without you.

Thanks, everyone!

DEDICATION

For Gary

Chapter 1

SOLAR NEUTRINOS

1.1 Neutrinos

Neutrinos, denoted ν, are electrically neutral, spin-$\frac{1}{2}$ particles. They were first proposed by Wolfgang Pauli in 1930 [31] to resolve the apparent nonconservation of energy in beta decays, as compared to the neatly energy-conserving alpha and gamma decays. An unseen particle in beta decays was deduced to carry away exactly the right amount of energy to clean up the discrepancy. Neutrinos enter the Standard Model as massless leptons, interacting only via the Weak force. However, observation of neutrino oscillations in experiments such as SNO [5], Super-K [1], and many others, provide compelling evidence that neutrinos in fact have a small but non-zero mass, making neutrinos one of the only areas of demonstrably Beyond-the-Standard-Model physics to date.

Neutrinos have been detected in three Weak eigenstates, called "flavors". Each flavor has a corresponding charged lepton with which it is associated: the electron, muon, and tau. The charged leptons are denoted e, μ, and τ, with the corresponding neutrinos denoted ν_e, ν_μ, and ν_τ. These associations are based on the observation that in charged current interactions a W boson always connects a charged lepton with its corresponding neutrino or antineutrino (for example, creating a μ and a $\bar{\nu}_\mu$, or destroying a μ and creating a ν_μ). There are also neutral current interactions, where a Z boson decays, for example, into a ν_τ and its antiparticle, a $\bar{\nu}_\tau$. In both charged current and neutral current interactions, not only is lepton number conserved, but flavor number is conserved as well. This means that if a reaction produces a ν_τ, it must also either destroy a τ or ν_τ, or it must produce a $\bar{\tau}$ or $\bar{\nu}_\tau$.

1.2 Neutrino Oscillations

While flavor number is conserved in Weak interactions, it is not conserved in general [33]. Once a neutrino is created (in some Weak eigenstate), it nearly always travels a considerable distance before being detected (in some Weak eigenstate) [14]. It was first observed by Super-K [1] and then confirmed by SNO [5] that the detected neutrinos are not necessarily the same flavors as when they were created. The results of multiple experiments point to a model for their transition known as neutrino oscillation. Because neutrinos nearly always travel extremely long distances between their creation and detection, one could in principle measure the mass of a neutrino with a time of flight experiment (by measuring its speed and energy, for example). Despite the fact that we do not yet actually have the capability to perform this measurement and detect a deviation from light speed [7], the fact that it could be done demonstrates why it must be the mass eigenstates, rather than flavor eigenstates, that govern propagation.

Imagine that a neutrino is created in an electron flavor eigenstate. This flavor eigenstate is a superposition of mass eigenstates, $|\nu_e\rangle = \sum_i U_{ei} |\nu_i\rangle$. Each mass eigenstate will develop a slightly different phase as the neutrino travels, where the phase is given by $\phi = Et - \vec{p} \cdot \vec{x}$. Expanding this, assuming the neutrino is hyperrelativistic, and using units where $c = 1$, results in $\phi = EL - \sqrt{E^2 - m^2}L$, or approximately $\phi = EL - EL\sqrt{1 - \left(\frac{m}{E}\right)^2}$. Because the neutrino mass is typically small compared to its energy, $\phi \approx \frac{m^2 L}{2E}$. The difference between the phases of two mass eigenstates will then be $\Delta\phi \approx \frac{\Delta(m^2)L}{2E}$. This phase difference results in a probability of detecting each flavor of neutrino that oscillates as a function of position and energy.

Oscillations in neutrinos traveling through matter are complicated by the fact that matter contains a preponderance of electrons but no positrons, so charge symmetry is broken by matter. A particularly interesting case of this is the MSW effect, where neutrinos produced in the core of the Sun transition adiabatically to the same mass eigenstate as the density of the solar material decreases at high radii. [41]

1.3 Solar neutrinos

The Sun and other stars produce nearly all their energy through fusion of hydrogen nuclei (protons) into helium nuclei [15]. The overall reaction can be written as $4p \to {}^4\mathrm{He}^{+2} + 2e^+ + 2\nu_e$. A helium nucleus contains two protons and two neutrons, so the positrons (e^+) are created to conserve charge, and the neutrinos are created to conserve lepton number. Because the products of this reaction are less massive than four protons, a net energy of 26.5 MeV is released each time this reaction occurs. Since we know how much energy the Sun radiates, we can calculate with small uncertainty the total rate of neutrino production in the Sun. For example, the irradiance of the Sun at Earth is about 1000 W/m^2, the surface area of a sphere with a radius matching Earth's orbit is about 3×10^{23} m^2, so the total power of the Sun is about 3×10^{26} W, which is about 2×10^{39} MeV/s, or about 2×10^{38} ν/s.

It turns out there are multiple pathways by which fusion occurs in stars. These are broadly divided into the "*pp*-chain" and "CNO cyle". The *pp*-chain is a series of reactions in which protons are built into deuterons, which then absorb another proton to become ^{3}He,[1] and then through various reactions produce ^{4}He. The CNO cycle starts with a ^{12}C nucleus which absorbs a series of four protons, undergoes two beta decays and emits a ^{4}He nucleus to return to ^{12}C.

In the Sun and smaller stars, the *pp*-chain accounts for most of the helium production. First, two protons combine to form a deuteron (a bound state of a proton and a neutron) via one of two processes:

$$\begin{gathered} p + p \to d + e^+ + \nu_e \\ \text{or} \\ p + p + e^- \to d + \nu_e \end{gathered} \tag{1.1}$$

The neutrinos from the first reaction are called "*pp* neutrinos" and constitute most solar neutrinos. These are the lowest energy solar neutrinos. The second reaction occurs at a much lower rate because it is a 3-body interaction, but its neutrinos (called "*pep* neutrinos")

---

[1] It is customary in nuclear physics to use an element's symbol to denote a fully ionized nucleus.

are monoenergetic, which distinguishes them from most other neutrino sources. The slow production of deuterons is the limiting step in fusion in the Sun. All subsequent steps occur on timescales of days or less, but the lifetime of the Sun is estimated to be about 10 billion years.

All deuterons then undergo the same next step: $d + p \to {}^3\mathrm{He} + \gamma$. The ^{3}He can then undergo three possible reactions. First,

$$ {}^3\mathrm{He} + p \to {}^4\mathrm{He} + e^+ + \nu_e \tag{1.2} $$

which accounts for only 10^{-7} of ^{4}He produced [11]. However, as these "*hep* neutrinos" are the highest energy neutrinos produced in the Sun, they may still be discernible experimentally. The majority of this dissertation will be spent discussing *hep* neutrinos.

The second reaction the ^{3}He may undergo is

$$ {}^3\mathrm{He} + {}^3\mathrm{He} \to {}^4\mathrm{He} + 2p \tag{1.3} $$

which accounts for about 84% of solar ^{4}He production. This reaction does not produce any neutrinos beyond those from the production of the ^{3}He.

Finally, the remaining 16% of the time the ^{3}He fuses with a ^{4}He:

$$ {}^3\mathrm{He} + {}^4\mathrm{He} \to {}^7\mathrm{Be} + \gamma. \tag{1.4} $$

The remainder of the *pp*-chain concerns the fate of the ^{7}Be produced in this reaction. The ^{7}Be can follow two pathways. First,

$$ \begin{gathered} {}^7\mathrm{Be} + e^- \to {}^7\mathrm{Li} + \nu_e \\ {}^7\mathrm{Li} + p \to 2\,{}^4\mathrm{He}. \end{gathered} \tag{1.5} $$

The neutrinos produced in this reaction are known as "^{7}Be" neutrinos. ^{7}Be neutrinos are interesting primarily because they are monoenergetic (or technically bienergetic, since the ^{7}Be can decay to two different states of ^{7}Li).

The ^{7}Be can also undergo the following reactions.

$$\begin{aligned} {}^7\text{Be} + p &\rightarrow {}^8\text{B} + \gamma \\ {}^8\text{B} &\rightarrow 2{}^4\text{He} + e^+ + \nu_e \end{aligned} \tag{1.6}$$

The neutrinos produced in this reaction are known as "^{8}B" neutrinos. ^{8}B neutrinos are important for this analysis particularly because SNO was primarily sensitive to ^{8}B neutrinos and therefore they are the largest background for studying *hep* neutrinos.

The CNO cycle is the other overall pathway for fusion in stars, and is thought to be the dominant fusion mechanism in stars even slightly larger than the Sun. In the CNO cycle, existing ^{12}C acts as a catalyst for fusion. The main cycle follows these reactions:

$$\begin{aligned} p + {}^{12}\text{C} &\rightarrow {}^{13}\text{N} + \gamma \\ {}^{13}\text{N} &\rightarrow {}^{13}\text{C} + e^+ + \nu_e \\ p + {}^{13}\text{C} &\rightarrow {}^{14}\text{N} + \gamma \\ p + {}^{14}\text{N} &\rightarrow {}^{15}\text{O} + \gamma \\ {}^{15}\text{O} &\rightarrow {}^{15}\text{N} + e^+ + \nu_e \\ p + {}^{15}\text{N} &\rightarrow {}^{12}\text{C} + {}^4\text{He}. \end{aligned} \tag{1.7}$$

These reactions produce two neutrinos, referred to as "^{13}N neutrinos" and "^{15}O neutrinos", respectively. A minor branch of the CNO cycle occurs when the ^{15}N absorbs the proton without emitting a ^{4}He:

$$\begin{aligned} p + {}^{15}\text{N} &\rightarrow {}^{16}\text{O} + \gamma \\ p + {}^{16}\text{O} &\rightarrow {}^{17}\text{F} + \gamma \\ {}^{17}\text{F} &\rightarrow {}^{17}\text{O} + e^+ + \nu_e \\ p + {}^{17}\text{O} &\rightarrow {}^{14}\text{N} + {}^4\text{He}. \end{aligned} \tag{1.8}$$

This branch produces "^{17}F neutrinos", which have a spectrum nearly identical to that of ^{15}O neutrinos.

Each of the eight types of solar neutinos (*pp*, *pep*, ^{7}Be, *hep*, ^{8}B, ^{13}N, ^{15}O, and ^{17}F) has a different spectrum and a different flux. Predicted spectra are shown in Figure 1.1, which is reproduced from [11]. A few features of these spectra bear mention. One is that the ^{8}B and *hep* neutrinos have dramatically higher energy than other types, which is why Cherenkov experiments like SNO (see Chapter 2) which have an energy threshold of a few MeV are only sensitive to these types of neutrinos. The ^{8}B spectrum ends at about 15 MeV and the *hep* spectrum ends at about 19 MeV, so even though the flux of ^{8}B neutrinos is about 1000 times larger than that of *hep* neutrinos, there is hope of measuring a clear *hep* signal above the ^{8}B endpoint. As a general trend, it is also the case that the types of neutrinos with a lower endpoint energy have a higher flux, such that the CNO, *pep* and ^{7}Be neutrinos have higher flux and lower energy than ^{8}B, and *pp* neutrinos have a much higher flux and lower energy yet.

1.4 Standard Solar Models

Standard Solar Models (SSMs) are our best models of the evolution of the Sun, taking into account all of the information we have available [12]. This information includes measured neutrino data, helioseismology measurements, solar composition information, and reaction cross-section information. An SSM begins with a zero-age main sequence star which is undifferentiated and has some assumed composition and other parameters. The model is stepped through time obeying physical constraints like energy conservation and hydrostatic equilibrium, and is constrained to have a final state consistent with current observations of the Sun.

One of the most important inputs into an SSM is the solar composition. The Sun is about 75% hydrogen and 24% helium, by mass, with other elements (in astronomy called "metals") making up the remaining 1% (largely oxygen, carbon, neon, and iron). Despite constituting a small fraction of the Sun's mass, the metals, particularly carbon, oxygen, and iron, have a significant impact on the rate of fusion in the Sun. Carbon and oxygen are important because they directly catalyze the CNO cycle, so any changes to their concentration will correspondingly change the rate of CNO processes. This is not the end of the story, however. Changing the rate of CNO processes will correspondingly change the temperature, pressure,

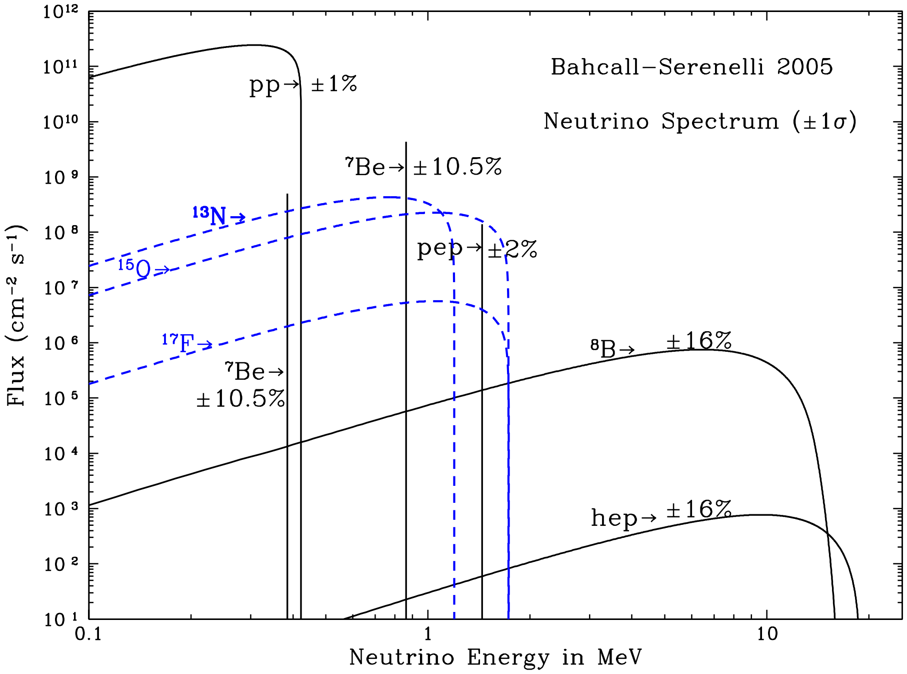

Figure 1.1: Predicted spectral fluxes of solar neutrinos. Note how the *hep* spectrum has the highest endpoint energy and the lowest flux of any of the solar neutrinos.

and density of the Sun, which will then have a feedback effect on all of the fusion processes.

The iron content of the Sun is important mainly because it has a dominant effect on the Sun's opacity, which is the ability for energy to radiate through the solar material. Since radiation is the dominant energy transfer mechanism out to about 73% of the Sun's radius (convection being dominant outside that radius), small variations in the concentration of iron will influence the opacity and change the temperature profile of the Sun which will also feed back into the reaction rates as discussed with the carbon and oxygen concentrations. Concentrations of other metals will also affect the solar opacity, though to a smaller degree.

The problem of calculating opacity under the conditions of the solar interior, even given elemental concentrations, is also difficult due to the variety of ionization states and inner electronic effects. [9]

Solar composition is determined by analyzing spectral properties of the photosphere, the visible layer of the Sun [8]. This is complicated by many things, including uncertainty of atomic (and molecular) spectra and particularly the fact that most metals make up very small fractions of the Sun's material. Another complication comes from the fact that we are viewing light that has passed through a considerable amount of solar atmosphere, so the results are also dependent on models of the solar atmosphere. Some allowance must also be made for the fact that the composition of the photosphere may differ from that of the inner layers of the Sun.

Determination of the solar content of rarer heavy elements is augmented by knowledge of meteorites. The five known samples of CI chondrites are considered to be the most pristine records of composition of the proto-solar material (the light gases having evaporated, leaving behind only the heavy metals). Mass spectroscopy of these meteorites produces a precise characterization of the elemental and isotopic starting conditions of the Sun.

Still, variations between methods and samples used produce different metallicity models, some of higher and some of lower metallicity [8] [32]. Uncertainty in solar metallicity is actually the dominant uncertainty in the theoretical rate of fusion in the Sun and therefore in the flux of solar neutrinos, for all types except *hep* neutrinos.

Another interesting result of SSMs is that the various fusion pathways are differentially preferred at various locations in the Sun, as shown in Figure 1.2, reproduced from [10]. We can make some sense of this preference by thinking about the *pp*-chain. First, we expect the rates of all reactions to drop off at high radius because the temperature, pressure, and density of solar material will all decrease. We also expect fewer neutrinos produced at low radius due to the purely volumetric consideration that there are fewer reactants in the smaller volumes at lower radii.

The *pp* neutrino curve gives us a baseline for the distribution of deuterons and therefore

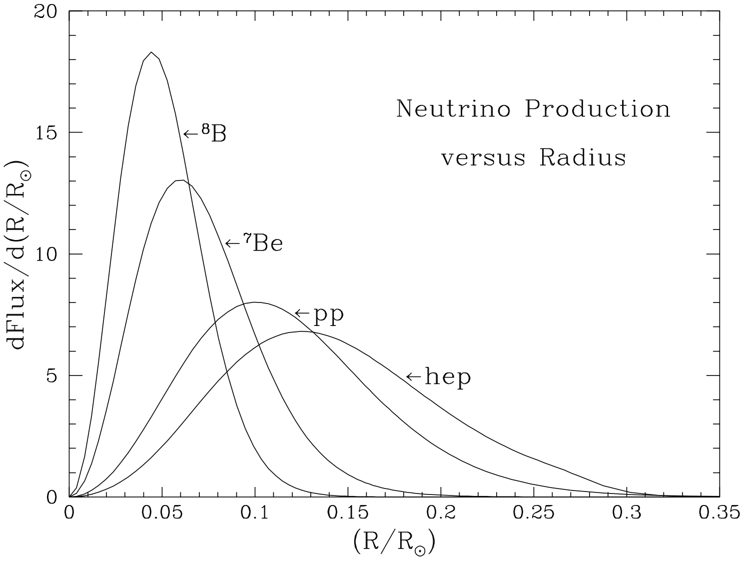

Figure 1.2: Probability distribution (as a fractional flux) of neutrino-producing fusion reactions as a function of radial position inside the Sun.

^{3}He nuclei, since these particles react much too quickly to diffuse far. The most preferred reaction for ^{3}He, with an 85% branching ratio, is $^{3}\mathrm{He} + {}^{3}\mathrm{He} \rightarrow {}^{4}\mathrm{He} + 2p$, which produces no neutrinos. Nearer the center of the Sun, there is a larger fraction of ^{4}He, which favors the production of ^{7}Be. The ^{7}Be picking up yet another proton to produce ^{8}B, instead of decaying to ^{7}Li, is most likely in the hotter, inner volume of the Sun.

The reason the *hep* reaction preferentially occurs at *higher* radius can be understood by this being the location where we expect ^{3}He to have the longest lifetime. At lower radius, other reactions dominate, but at higher radius the diffuse ^{3}He has more time to fuse with a proton instead of finding another ^{3}He or ^{4}He. Note however that the *hep* reaction still occurs at a far lower rate than the others, it's just that most of the *hep* reactions occur at higher radii than the other reactions.

1.5 hep neutrinos

Direct or indirect measurements of ^{8}B, ^{7}Be, *pep*, and *pp* neutrino rates have all been conducted. Borexino, for instance, has measured all of these (see, for example, [13]). Upcoming experiments, for example SNO+ [6], hope to detect CNO neutrinos. However, *hep* neutrinos are particularly elusive due to their very low flux.

Nevertheless, a measurement of *hep* neutrinos is interesting for a number of reasons beyond simple scientific curiosity. For one thing, the theoretical rate of *hep* neutrinos is the only solar neutrino rate whose uncertainty is not dominated by the uncertainty in solar metallicity, so it provides a somewhat independent test of the Standard Solar Model. The reason for this is partly because *hep* neutrino production is spread out over a larger volume of the Sun, so small changes to the solar temperature profile have a smaller effect on *hep* neutrinos than on the other processes. Instead, due to considerable cancellation between terms, the largest uncertainty in the rate is from the nuclear matrix element [30], and it seems likely that this could be improved if there was an experimental result to compare it to.

1.6 Solar neutrino experiments

Previous experiments relied on two types of interactions to detect solar neutrinos. One is called inverse beta decay, which the famous Homestake experiment [18] used to detect solar neutrinos: $\nu_e + {}^{37}\text{Cl} \rightarrow e^- + {}^{37}\text{Ar}$. Inverse beta decay is only sensitive to electron flavor neutrinos[2]. By counting decays of the resulting ^{37}Ar atoms, it is possible to determine the

---

[2] Actually, a muon neutrino can trigger an analogous reaction called a charged current quasi-elastic reaction (CCQE) that produces a muon instead of an electron. However, solar neutrinos do not have the required 104 MeV of energy to create a muon.

rate of argon formation, and therefore the rate of neutrino interactions with the chlorine.

The other type of reaction previously used to detect solar neutrinos is elastic scattering of electrons. Experiments such as Super-K [1] relied on this technique to detect neutrinos. In this process, a neutrino simply exchanges energy and momentum with an electron. If a neutrino elastically scatters off an electron in water and the electron gains enough energy, it will produce a Cherenkov cone that can be detected.

Cherenkov radiation occurs any time a charged particle moves faster than the speed of light in a given medium. In a way analogous to a sonic boom, wavefronts from the particle add coherently along a forward-pointing cone. The angle of the cone depends on the speed of the particle and the index of refraction of the medium. The direction of the cone gives the direction of the particle's travel, and the intensity gives information about the particle's energy.

Elastic scattering, unlike inverse beta decay, does have some sensitivity to ν_μs and ν_τs, but it still is much more sensitive to electron neutrinos (by about a factor of 6). For this reason, along with the different energy ranges of the two experiments, the solar neutrino deficit observed in Super-K was a factor of about one half, while Homestake observed about one third of the expected flux.

A nice feature of detecting a Cherenkov cone from elastic scattering events is that the direction of the electron is easily measured. Because both neutrinos and electrons have very small masses compared to the center-of-mass energies of these interactions, a significant portion of the neutrino's initial momentum will be transferred to the electron. Therefore, the direction of the electron will be well-correlated with that of the neutrino. An unhelpful consequence of the small masses of the neutrino and electron is that the neutrino and electron will share the kinetic energy about equally on average. This means that the energy of the electron will not be a faithful indicator of the energy of the neutrino, and indeed, a significant fraction of elastic scattering events may result in electrons without enough energy to be detected at all.

SNO (see Chapter 2) was able to detect elastic scattering interactions in the heavy water (and even the light water) by detecting the Cherenkov radiation from the electron. Additionally, neutrinos are able to undergo two more reactions when they interact with a deuterium nucleus, called charged current and neutral current interactions. In SNO, charged current interactions follow the reaction

$$\nu_e + d \rightarrow 2p + e^-, \tag{1.9}$$

where d denotes a deuteron. For the purposes of SNO, charged current interactions are only sensitive to ν_es because, as with the inverse beta decay discussion above, solar neutrino energies are much smaller than a muon's rest mass. In fact, this reaction basically is an inverse beta decay, except the "daughter nucleus" is two unbound protons.

Because the protons have very large masses compared to the electron, kinematics requires that in a charged current interaction the electron will come away with nearly all of the neutrino's kinetic energy (minus that lost in the inelastic collision). Because of this, the electron's final state energy in a charged current interaction is more strongly correlated with the neutrino's initial energy than in an elastic scattering interaction. Because the final state of a charged current interaction is a three-body state, the momentum is distributed among the protons and electron, so the electron final state momentum is almost totally uncorrelated to the incoming neutrino direction. It turns out, in fact, there is a slight *anti*-correlation between the neutrino's and electron's momenta. This can be understood as arising from conservation of angular momentum. Deuterons are spin-1, but the final state includes two protons, which will preferentially have anti-aligned spins. In order for the proton spins to end up anti-aligned, the neutrino-deuteron system must have had total spin of $\frac{1}{2}$, in the direction opposite to that of the neutrino's spin. Therefore, the final state left-handed electron will be preferentially emitted in the direction opposite to the neutrino's momentum. Unfortunately, this effect is too small to be helpful for this analysis.

A neutrino interacting with a deuterium nucleus can also undergo the neutral current interaction:

$$\nu + d \rightarrow p + n + \nu. \tag{1.10}$$

The neutral current interaction is equally sensitive to all neutrino flavors because no Feynman diagrams contribute where a particle changes type. Unfortunately, since no electron is produced in this reaction, it is impossible to learn anything about the incident neutrino's energy, and therefore these events won't be useful for the analysis in this work. However, detection of the neutron provides clear evidence of the neutral current reaction (assuming sufficiently limited background neutrons), and was critical for SNO's neutrino oscillation result.

SNO was conducted in three phases which differed in their strategies to detect the neutral current reaction neutrons. In Phase I, the neutrons were captured on D_2O, emitting a 6.25 MeV γ-ray [5] which was then detected via Cherenkov radiation of subsequent Compton electrons and e^+e^- pairs. In Phase II, 2000 kg of NaCl was added to the D_2O. Because ^{35}Cl has a much larger neutron capture cross-section than D_2O, this resulted in an increase in detection efficiency. The $n + {}^{35}\mathrm{Cl} \rightarrow {}^{36}\mathrm{Cl} + \gamma$ reaction produces several γ-rays totalling 8.6 MeV. In Phase III, an array of neutron counters SNO called Neutral Current Detectors were installed in the detector. These are described further in Chapter 2.

No experiment has yet detected *hep* neutrinos because of their low flux. Nevertheless, their detection is possible because they extend to a higher energy than any other solar neutrino (about 18.8 MeV vs. 15 MeV for ^{8}B neutrinos). This is also higher than nearly all of the radioactive backgrounds. The SNO experiment (see Chapter 2) has a few advantages over any other experiments for detecting *hep* neutrinos. For one thing, it is sensitive to neutrino events in the ^{8}B and *hep* regimes, but not overwhelmed by the large flux of solar neutrinos at lower energies. Additionally, SNO has an advantage over larger water-Cherenkov detectors like Super-K because light water only permits elastic scattering events, while heavy water also permits charged current events. In both types of events, the detector measures the energy of an outgoing electron, but an important difference is that, as described above, charged current reactions are much better at preserving the incoming neutrino's energy, which is the only handle to discriminate *hep* from ^{8}B events. Since no new experiments are planning to use D_2O as a target, SNO data may be the best bet for ever detecting *hep* neutrinos.

SNO data has previously been analyzed in an attempt to find *hep* neutrinos [2]. That work used a box analysis technique which found a number of events consistent with the expected background rate, setting a limit of 2.3×10^4 cm^{-2}s^{-1} at 90% confidence level. That work only used data from the first of the three SNO phases. For comparison, the predicted rate is $7.97 \pm 1.24 \times 10^3$ cm^{-2}s^{-1} [12]. A second analysis [23] made use of a spectral analysis and used three phases of data. This analysis gave tantalizing hints of a signal with a flux of $2.40^{+1.19}_{-1.60} \times 10^4$ cm^{-2}s^{-1}, albeit at less than 2σ significance.

In this work, we expect to be able to improve the sensitivity to *hep* neutrinos by increasing the fiducial volume, including additional runs that were not used in previous analyses, and using an improved event fitter that returns a posterior energy distribution for each event rather than just a maximum likelihood value.

Chapter 2

THE SUDBURY NEUTRINO OBSERVATORY

The Sudbury Neutrino Observatory (SNO) was a neutrino detector whose primary purpose was to measure all flavors of neutrinos arriving from the sun. This was motivated by an apparent deficit in the total flux of solar neutrinos measured by several previous experiments relative to fairly strict theoretical predictions. This deficit, known as the Solar Neutrino Problem, was conjectured to be due to a phenomenon known as neutrino oscillation, wherein a neutrino may be detected in a different Weak eigenstate than the one with which it was created. Since earlier experiments had little or no sensitivity to ν_μs or ν_τs, the neutrino oscillation hypothesis explained the apparent deficit by supposing that the predictions for total neutrino flux were correct, but some portion of the neutrinos had oscillated to invisible flavors. In order to test this hypothesis, SNO was designed to be able to detect multiple types of interactions, one of which (the charged current interaction) is only sensitive to electron-flavor neutrinos, while another (the neutral current interaction) is equally sensitive to all flavors, and in this way the neutrino oscillation hypothesis (that we receive the expected flux of neutrinos, but with unexpected flavors) could be tested.

2.1 SNO detector

The SNO detector was located about 2 km deep in an active nickel mine in Sudbury, Ontario, Canada, which makes it the deepest neutrino experiment to date. The rock overburden (6000 meters water equivalent) reduces the cosmic muon flux dramatically. Despite residing in a mine, the area around the detector is treated as a clean room in order to minimize radioactive contamination of the detector.

The detector itself (see Figure 2.1) consists of just under 10000 Hamamatsu photomul-

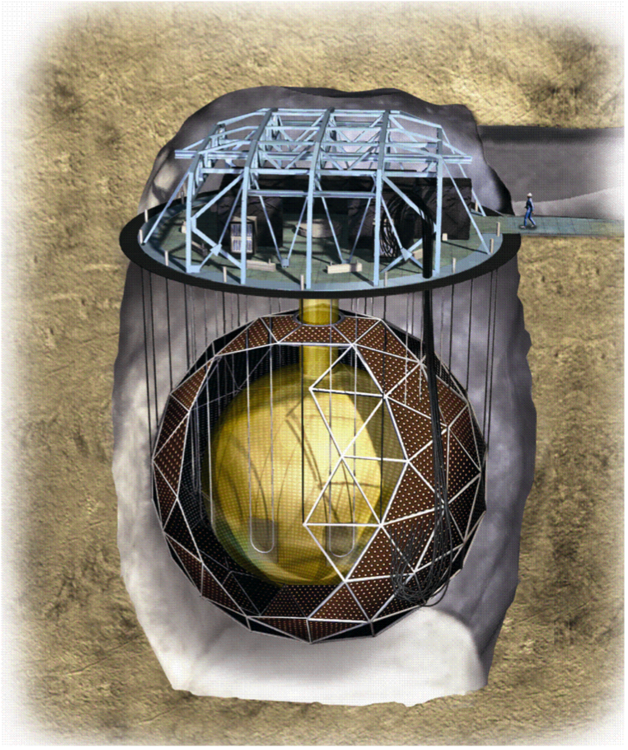

Figure 2.1: SNO detector. The AV is in yellow, the PSUP is in brown, the cavity is in gray, and the deck is in blue. A person is shown for scale.

tiplier tubes (PMTs) pointing inward and arranged on an 8.24 m-radius sphere called the PSUP. To increase the light collection, each PMT is surrounded by a ring of aluminum petals collectively called "concentrators". The concentrators are arranged in such a way that they redirect light from the central part of the detector into the PMTs, but do not do so for light that approaches the PMTs from too large an angle (or at least do so less efficiently). The concentrators exhibited dramatic and uneven aging over the course of data collection, with some concentrators corroding and others falling off completely, which caused decreased light collection and also increased PMT-to-PMT variation over time. Additionally, over the course of the experiment, roughly 10% of the PMTs failed for various reasons.

Concentric with the PSUP is the "acrylic vessel" (AV), a 5.5 cm-thick, 6 m-radius acrylic spherical shell with a cylindrical “neck” extending upwards through the PSUP. The PSUP is in a cavity filled with ultrapure water, and the PSUP itself is also filled with ultrapure water. The AV is filled with D_2O where D denotes deuterium, an isotope of hydrogen with one proton and one neutron. As normal water has molecular weight 18, but D_2O has molecular weight 20, D_2O is 11% more dense than normal water and is therefore called “heavy water”. In order to support the weight of the D_2O, an array of “hold-up ropes” run between the AV and the support structure above the detector. The hold-up ropes suspend the AV by wrapping around acrylic plates called belly plates, which are attached at the equator of the AV.

While most of the PMTs point inward toward the AV, a small number of tubes are outward-looking (“OWL”), and act as a muon veto. Another small number of the inward-looking tubes are low-gain tubes, which are wired such that in the event of a large amount of light in the detector these tubes will still operate normally. There are also two tubes specially placed in the neck of the detector to help veto events occurring in this geometrically inconvenient area of the detector.

Set into the rock surrounding the detector is a series of 14 “compensation coils” which cancel out the Earth’s magnetic field. These are important because PMTs operate by accelerating electrons over a high voltage (around 2000 V, though this is tuned for each tube), and stray magnetic fields substantially decrease their ability to do so.

Above the detector is a “deck” which houses racks of electronics to read out and control the PMTs. In the middle of the deck is an access hatch to the neck of the acrylic vessel. This allows, among other things, calibration sources to be inserted into the detector. Sources can be moved around by a series of ropes called the “manipulator”.

The SNO experiment ran in three phases, which employed different methods for detecting neutrons. In Phase I, the detector was arranged as described above and neutrons were detected by capturing on deuterium. In Phase II, the only change was that NaCl was added to the D_2O region in order to allow neutrons to capture on ^{35}Cl. In Phase III the detector was upgraded with a series of 40 nickel tubes installed inside the AV. Of these, 36 were proportional counters filled with ^{3}He that capture neutrons via the reaction $^3\mathrm{He} + n \rightarrow {}^3\mathrm{H} + p$. The protons and tritons move through the helium gas creating ionization electrons which produce detectable voltages in the tubes. Because the proportional counters (hereafter referred to as "NCDs" or "Neutral Current Detectors") did not rely on Cherenkov radiation from neutron capture γs, they provided a nice cross-check on the rate of neutral current reactions. In addition to the 36 NCDs, there were four tubes containing ^{4}He to test for alpha backgrounds within the tubes themselves.

2.2 Operation

Inside the SNO detector, a variety of physics events could occur which would produce photons. These photons traveled through the detector and some fraction struck the active surface of a PMT. A photon that strikes a PMT may produce a photoelectron which then accelerates across a high voltage difference to a dynode. Upon striking a dynode, the electron "multiplies" by releasing several electrons, which then each accelerate toward additional dynodes. After electrons cascade down a series of dynodes, a measurable current is detected at the anode.

The SNO detector was monitored electronically for close time coincidences between many PMT hits. A variety of triggers were used to help ensure that interesting events were not missed. Upon receipt of a trigger, the charge, time, and some other information about the PMT hits were recorded. Software called an event builder continuously partitioned PMT hits into "events" (based on triggers) to facilitate further analysis.

Detector operators were on shift monitoring the overall state of the detector. They were able to address unusual behavior through a variety of channels, up to shutting down the detector. Operators were responsible for keeping shift reports and making sure operating procedures were followed so that data was not compromised.

Once the data was partitioned into events, a series of low-level tests called Data Analysis Mask Numbers (DAMN) were applied to the events. The DAMN IDs are useful for identifying certain backgrounds (such as muons and flashing PMTs), events that are troublesome for analysis (such as those in the neck of the AV), events that look like electronics noise, and events in close time coincidence. By cutting on the DAMN ID, analyses can choose which types of events to analyze. Events would then be fit for position and energy by whichever fitter was desired, generally based on the number and pattern of photons detected, and analysis would proceed. For more on fitters, see Chapters 3 and 4.

The SNO detector was also subjected to a variety of calibrations. Some calibrations ensured that the timing of PMTs was consistent. Some calibrated the PMT and NCD electronics. Others were radioactive sources used to calibrate the energy and position reconstruction of electrons, neutrons, and radioactive background events.

A number of radioactive backgrounds were present in SNO [17]. Because the detector was located in a mine, elevated radon levels left radioactive daughters on any surface that was exposed to mine air, and any dust contaminating the detector was also radioactive. Additionally, the borosilicate glass of the PMTs is radioactive, which is part of the reason for the two meters of H_2O between the AV and the PSUP. Furthermore, it is nearly impossible to remove the trace amounts of radioactive uranium and thorium that are present in nearly everything. Finally, γs of modest energy are able to dissociate deuterons, which mimics neutral current events. This concern dominated a lot of SNO's radiopurity considerations.

Besides radioactive contamination, cosmic rays produce events in the detector. High-energy particles from space, probably protons, strike atoms in the atmosphere and produce showers of pions, which decay to muons. While the vast majority of these muons do not make it to the depth of the SNO experiment, a small flux did reach the detector, amounting to 0.27 $\mu/\mathrm{m}^2/\mathrm{d}$ [4]. Muons themselves are pretty easy to identify, as they produce a lot of light, including outside PSUP. However, they can create cosmogenic backgrounds in the water and acrylic which decay over longer time scales (specifically, spallation of $^{16}\mathrm{O}$ nuclei can produce nearly any nucleus with mass up to 16, and there is also a possibility of heavier daughter nuclei from $^{17}\mathrm{O}$ and $^{18}\mathrm{O}$). Of particular concern for the *hep* analysis is the isotope $^{12}\mathrm{C}^*$, which can decay by releasing a γ at 15.1 MeV [20]. This isotope is produced at an unknown but very small rate when a muon strikes an $^{16}\mathrm{O}$, causing it to release a neutron or proton to become $^{15}O^*$ or $^{15}N^*$, which can then cascade down to $^{12}C^*$.

The detector itself also was responsible for a few types of backgrounds. Bubbles, PMTs, and external sounds from the mine (e.g. blasting and "rock bursts", which are minor collapses, happen audibly multiple times a day) all produced events in the detector. From time to time, the acrylic vessel also produced isotropic flashes of light whose source is not well-understood. These will be addressed specifically in Chapter 3.

2.3 SNOMAN

Complementary to the SNO detector is a simulation software called SNOMAN (SNO Monte carlo ANalysis code) [37]. SNOMAN uses a suite of physics engines including GEANT, EGS4, NUANCE, MUSIC, and many others, along with detailed information about the detector and customizable control routines to produce simulated ("Monte Carlo") events under specific conditions. By simulating signal and background events, the Monte Carlo data can then be used to inform analysis on the real data, such as testing and refining event fitters, effects of systematics on measurable results, testing hypotheses about the detector, and testing calibration and signal extraction.

SNOMAN was developed by dozens of members of the SNO collaboration over many years, and its validity was tested rigorously. Still, the code is limited by our limited knowledge, including empirical inputs, uncertainty in cross-sections, particle fluxes, and variations in detector components spatially and temporally. Even in these cases, care was taken to incorporate the results of in situ measurements and to cross-check with data-driven methods to whatever degree possible.

2.4 SNO+

Although SNO is complete, the detector is being refurbished for another experiment, called SNO+. The detector is largely the same, though the electronics have been upgraded to allow for much higher data rates, and the D_2O is being replaced with a liquid scintillator called linear alkyl benzene (LAB). With liquid scintillator, the energy threshold for elastic scattering events will be much lower, and will perhaps allow for detection of CNO neutrinos, *pp* neutrinos, and *pep* neutrinos, if backgrounds can be sufficiently reduced. It will also check KamLAND's measurement of Δm^2_{21} [34]. It is also possible that a continental-crust geoneutrino measurement can be made at levels that are useful for geophysics (a particularly appealing possibility given the extremely well-studied geology of the Sudbury Basin). The primary purpose of SNO+, however, is to load the LAB with ^{130}Te, a double-beta decay isotope, in order to search for neutrinoless double-beta decay. The observation of neutrinoless double-beta decay would have important consequences for cosmology and particle physics because it would prove neutrinos are their own anti-particles, a possibility which is suggested by certain theories that leptogenesis may have contributed to the matter-antimatter imbalance in the universe. SNO+ has a number of advantages over other neutrinoless double-beta decay searches. One major advantage is that the detector design is not predicated on the details of the element used, as it is in single-crystal germanium detectors, liquid xenon detectors, and tellurium oxide crystal bolometers, so if neutrinoless double-beta decay is observed, SNO+ could do a cross-check with other isotopes with minimal additional development. Another benefit is the ability to deploy a large mass of double-beta decay isotope relatively easily.

Chapter 3

EVENT FITTER

In this chapter, we describe the fitter for single events. This fitter is based on a Bayesian analysis, which means that we use Bayes's Theorem to calculate a "posterior" probability distribution for several parameters of interest for each event. For example, we can calculate $P(E|\text{data})$, the probability the event had some energy E given the measured data. Bayes's Theorem requires a "likelihood function" and a "prior", detailed below. The likelihood function is a measure of how closely the data match expectations based on given parameters. In our case, the data consist of which PMTs detect photons, and the times at which they do so. The expected numbers and arrival times of photons at each PMT are based on spatial and temporal probability density functions built from a variety of parameters of the particle generating the light such as energy, position where an interaction occurred, and direction of travel. The posterior distributions are determined using a Markov Chain Monte Carlo algorithm.

3.1 Bayesian Basics

Bayes's Theorem states that for probabilities P,

$$P(\Psi|\Xi) = \frac{P(\Xi|\Psi)P(\Psi)}{P(\Xi)} \tag{3.1}$$

where the function inputs before the | are taken to be independent variables and the inputs after the | are taken to be given parameters. This theorem is simply a result of conditional probabilities; one need only rearrange the second equality in $P(\Psi, \Xi) = P(\Psi|\Xi)P(\Xi) = P(\Xi|\Psi)P(\Psi)$ to derive it. In a Bayesian analysis, the really useful part is assigning to Ξ our *data*, meaning measured quantities (PMT positions, photon arrival times, etc.), and assigning to Ψ *parameters*, meaning quantities we wish to derive (energy, event type, etc.).

The various parts of Bayes's Theorem are given names. $P(\Psi|\Xi)$ is called the "posterior" and we interpret it as the probability density of an underlying set of parameters given all the information we have available. $P(\Xi|\Psi)$ is called the "likelihood function". This is a manifestly calculable function since one assumes a given distribution of parameters and then calculates the probability (density) of getting the actual data. I will refer to the quantity $P(\Xi)$ as the "normalization" because it serves to normalize the likelihood function so that the posterior will satisfy all the rules of probability. In simple cases, $P(\Xi)$ can be calculated by writing

$$P(\Xi) = \int_{\text{all } \Psi} P(\Xi|\Psi) P(\Psi) \, \mathrm{d}\Psi, \tag{3.2}$$

but in our case we will not need to calculate it explicitly as normalization will come directly from the Markov Chain Monte Carlo algorithm as discussed in section 3.2.

The quantity $P(\Psi)$ is called the "prior" and is interpreted as expressing all information we have about the parameters Ψ that may not be contained in the data Ξ. For example, previous experiments may tell us information about a quantity, but with some uncertainty, or physical constraints may not allow certain quantities to be negative. We very often have some expectations about what the parameters should be, and the prior allows us to incorporate those (possibly subjective) expectations in a transparent way. If the data are good enough, the choice of prior should not influence the result strongly.

Perhaps a simple example is in order to demonstrate these ideas. Let us assume you wish to test whether a coin toss really results in a 50% chance of coming up heads. You do an experiment where you flip the coin 10 times and record 8 heads. In this case, Ψ is the real rate of getting heads (which is what you want to know), and Ξ is the rate you measure (8/10). You can calculate the probability of getting 8 heads in 10 tosses given a fair coin ($\Psi = 0.5$) using the binomial formula: $P(\Xi = 8/10|\Psi = 0.5) = 0.044$. This is the likelihood factor in Bayes's Theorem. Similarly, the binomial formula can be used to find $P(\Xi = 8/10|\Psi = 0.8) = 0.30$. Explicit calculation of $P(\Xi = 8/10)$ as in Equation 3.2 reveals $P(\Xi = 8/10) = \frac{1}{11}$ (which makes sense since there are 11 possible outcomes from

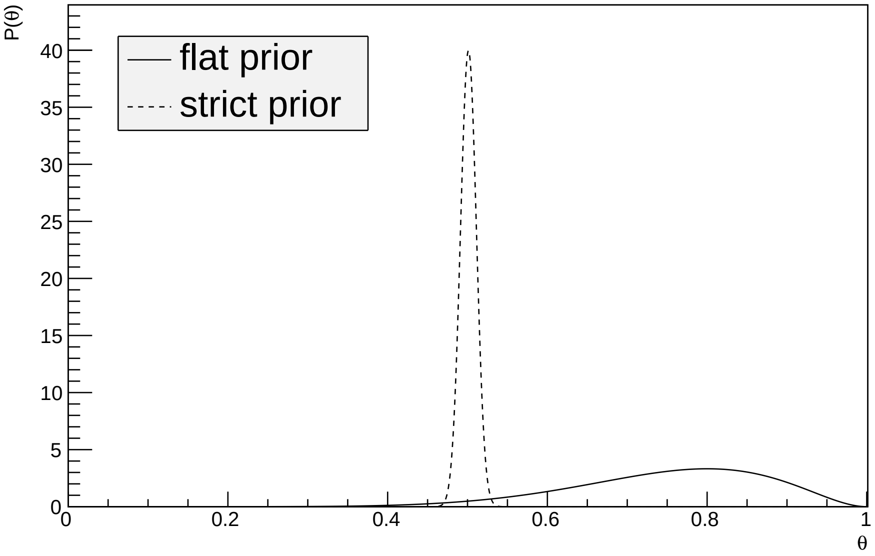

Figure 3.1: Posterior probability for true rate of flipping heads on a coin that has produced 8 heads in 10 trials. The solid line assumes a flat prior. The dotted line assumes a Gaussian prior centered on 0.5 with a width of 0.01. In this case, the prior has a strong influence on the posterior, so the test does not convincingly measure the true rate of flipping heads.

flipping a coin 10 times). If you chose a "flat prior", meaning all values of Ψ are equally likely, you would happily conclude that the coin probably comes up heads about 80% of the time (though you wouldn't rule out other values of Ψ since the posterior is quite broad). The solid line in Figure 3.1 is the posterior with a flat prior. However, my experience is that coins are generally pretty fair, and based on my prior distribution (perhaps a Gaussian of width 0.01 centered on 0.50, with additional peaks at 0.00 and 1.00 because this may be a trick coin), the posterior distribution would still be peaked at 0.5. The dashed line in Figure 3.1 is the posterior distribution with my more strict prior. Since the prior strongly influences the posterior in this case, this experiment does not convincingly measure Ψ. However, if you built a machine to flip the coin 1 million times, and it came up heads 520,000 times, this would be *very* convincing evidence that Ψ is close to 0.52. With any reasonable prior (flat, peaked, etc.) the posterior distribution would be very sharply peaked near 0.52 because the likelihood function is so sharply peaked at this value.

3.2 Markov Chain Monte Carlo

In practice, one does not necessarily need to calculate the likelihood, prior, and normalization explicitly to extract the posterior probability. Using a Markov Chain Monte Carlo (MCMC) technique one can create posterior distributions for each parameter without ever finding the normalization.

The general idea behind an MCMC is to draw a long series of samples from the parameter space in such a way that the distribution of these samples mirrors the underlying probability distribution for the parameter. First, we evaluate the likelihood[1] at a starting point in parameter space, then choose a random but nearby point. If the likelihood at the nearby point is "good" based on some criteria, we accept that point and then choose our third point near to this second point. If, however, the second point had a "bad" likelihood, we would reuse the original point as our second step, and choose our third step from near this original point. As the algorithm steps through many points in parameter space, it preferentially spends time where the likelihood is "good". The output of the MCMC is the distribution of parameters over all the steps (or some subset of the steps, as discussed later), so normalizing the histograms of each parameter results in correctly normalized posterior distributions without ever explicitly calculating normalization factors.

To be specific about what constitutes a "good" likelihood, we use a Metropolis-Hastings algorithm. In this algorithm, we always accept points with better likelihoods than the previous step, and sometimes accept points with worse likelihoods. If the likelihood is worse, the probability the point will be accepted is proportional to the ratio of the new to old likelihood. For example, if the old point has a likelihood of 0.1, and we sample a new point with a likelihood of 0.2, we always accept the new point. Now the point with a likelihood

---

[1] Actually, we will calculate the likelihood times the prior, $P(\Xi|\Psi)P(\Psi)$. In the simplest case the priors are flat, and this reduces to the likelihood.

of 0.2 becomes the “old” point, and we choose a new point near this point. If our new point has a likelihood of 0.1, there is a 50% chance of accepting the new point, and a 50% chance of using the old point again. In this way, more time is spent sampling near the best points, but the algorithm will eventually sample the whole parameter space with a probability proportional to the likelihood.

A choice one must make when using an MCMC is how to choose the new point near the previous point. This choice is important, because if the steps are too small, it will take too many steps to sample the whole parameter space. However, if the steps are too large, then nearly all the steps will take the algorithm too far away from “good” likelihoods, and will be rejected (which, again, will fail to sample the whole parameter space). When dealing with steps in very different spaces (for example, spatial coordinates, directions, energy, and fractions), it is hard to know if steps in one parameter are dominating whether jumps are accepted or rejected. To be systematic about choosing step sizes, one can calculate autocorrelations for the parameters. The autocorrelation of a set of data x is its correlation with itself as a function of offset τ

$$A(\tau) = \frac{E[(x_i - \mu)(x_{i+\tau} - \mu)]}{\sigma^2}, \tag{3.3}$$

where E is the expectation value and μ and σ are the mean and standard deviation of x. High autocorrelations at long offsets indicate step sizes that are too small (because the parameter barely changes from step to step). The indication that step sizes are too large is that too many steps are rejected. The rule of thumb is that about 25% of the steps should be accepted, but unfortunately failing this rule of thumb does not tell you which parameters have too large of step sizes.

The distribution from which new steps are drawn also must be decided. Typically one uses Gaussian jumps, though other choices are fine as well. We use Gaussian jumps in most parameters. There is, however, some subtlety in this choice. For example, if three fractions must add up to 1, making Gaussian jumps in two of them and then choosing the third so they add to 1 will not result in equally-sized Gaussian jumps in the third fraction. On the other hand, making Gaussian jumps in all three and then normalizing them to 1 will not result in Gaussian jumps for any of them. Again, this choice should not be important, but for completeness I mention that we choose the second method in these cases. Another case where this choice is tricky is when choosing a new direction near an old direction. Because Gaussian distributions have infinitely long tails, in principle there is some wraparound (if an angle greater than π radians is chosen) which could result in unexpected distributions. We take a computationally simple approach and choose small-angular Gaussian jumps around the x-, y-, and z-axes and perform the small-angle rotations simultaneously.

Incorporating priors into an MCMC is straightforward. One simply chooses a distribution for the prior, and multiplies the likelihood by its value with the given parameters at each step. This way, the algorithm accepts points based on the product of the likelihood with the prior, rather than just the likelihood, but normalization is still achieved automatically as before.

A drawback of using an MCMC is that each step is necessarily highly correlated with the steps before and after it. One can solve this by only using every nth step, such as every 1000th step. Usually one doesn't need to worry much about this correlation because one is just interested in the overall distribution of parameter values, rather than series of nearby steps. One should be aware, however, that these distributions typically fail chi-square tests because of the propensity for the MCMC to remain in the same spot for potentially hundreds of steps, creating a spiky distribution. A second drawback of MCMC analyses is that they are computationally intensive, which is particularly true if only a small fraction of the steps are used. Related to the computational intensity is the need to throw away a portion of the steps near the beginning that correspond to steps with bad likelihood before the MCMC finds any good points at all (a period called "burn-in").

A nice feature of MCMCs is that it is easy to see if certain variables are correlated with each other. All one needs to do is calculate the correlations of the parameters over the series of steps.

3.3 Physical model of events

In order to fit events, we developed a physical model for the expected distributions of photons. Our model for the most common type of event contains three sources of photons: Cherenkov radiation, reflections or scattering, and backgrounds.

The main source of photons is electrons that produce Cherenkov radiation by traveling faster than the speed of light in their medium. These electrons are assumed to have been created by neutrino charged current and elastic scattering events in the heavy water, so they provide information about the incoming neutrinos (See Chapters 1 and 2). Some of the Cherenkov photons are reflected off surfaces or scattered within media and may arrive anywhere in the detector at times somewhat later than the prompt Cherenkov photons. The last source of photons (actually PMT hits, since they are not necessarily photons) we model is what we call "background". The background corresponds to PMTs firing randomly in a way unrelated to a real event, which all PMTs do at some characteristic rate called the "dark rate".

3.3.1 Cherenkov events

When a charged particle travels faster than the speed of light in a medium, it creates wavefronts that add coherently along a cone at an angle $\theta_c = \frac{1}{n\beta}$ radians from the path of the particle. Here, n is the index of refraction of the medium and β is the speed of the particle, as a fraction of the speed of light in a vacuum. For D_2O, n is frequency-dependent (and very similar to H_2O). In the wavelength range we are interested, the index averages 1.337. For any events we consider in SNO, $1 - \beta \ll 0.01$.

Ideally, these "Cherenkov photons" are produced in an infinitely narrow cone. Effects such as energy loss and finite path lengths between tiny scatters off nuclei and electrons in the medium will give the cone some width [25]. Even with these considerations, for most particles (such as muons), the cone remains exceedingly sharp. However, electrons are so light that as they scatter multiple times in the water the Cherenkov cone becomes quite

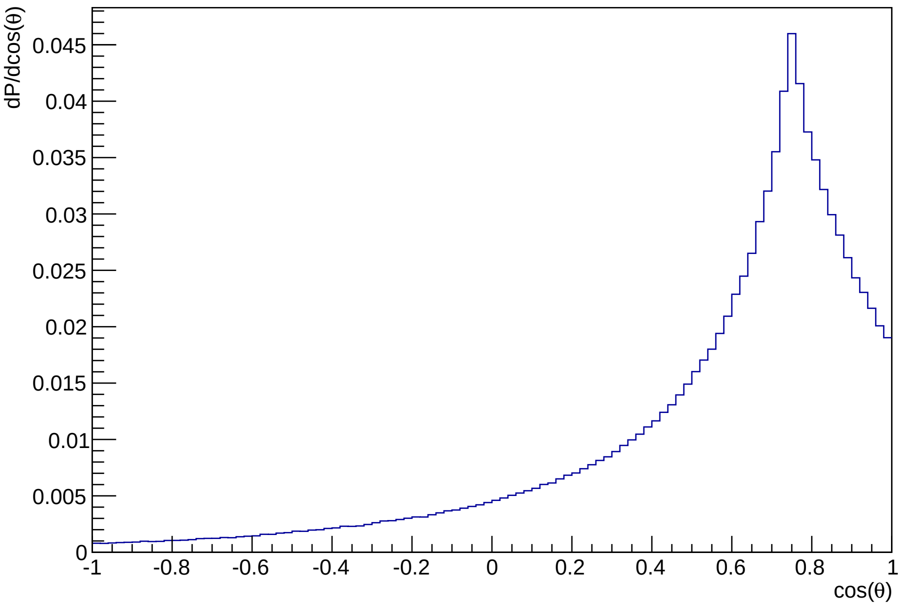

Figure 3.2: Distribution of Cherenkov photons from an electron in water, relative to the electron's initial direction.

blurred. A discussion of multiple scattering can be found in [16]. For this reason, the spatial distribution for Cherenkov photons takes on the peculiar shape in Figure 3.2. This histogram was generated by the SNOMAN Monte Carlo. See Chapter 2 for a discussion of SNOMAN.

An improvement that this work makes over previous Cherenkov cone analyses is to allow for multiple Cherenkov cones. For the same reasons that electrons produce broadened cones, they also have the possibility of undergoing a fairly large scatter and producing a second cone in a different direction. In fact, an electron is much more likely to undergo a single large scatter than a succession of smaller scatters that cumulatively produce a large deflection (see [16]). The distribution shown in Figure 3.2 is an *average* of many events, but the distribution in individual events may differ quite radically from this average, particularly if the electron undergoes a large scatter.

Our approach then is to fit multiple cones with some other (narrower) distribution for each cone. Each cone, k, is azimuthally symmetric about its axis, so it will have an individual distribution that is only a function of θ_k, the angle from the axis of cone k. Each cone will also have a relative weighting f_k, giving an overall event distribution

$$\frac{\mathrm{d}P_{\text{Cherenkov}}}{\mathrm{d}\Omega}(\theta, \phi) = \frac{1}{2\pi} \sum_k f_k \frac{\mathrm{d}P_k}{\mathrm{d}\cos(\theta_k)}. \tag{3.4}$$

The $\frac{1}{2\pi}$ factor gives the correct normalization since each individual cone is symmetric in ϕ_k.

Spatial distribution

In order to derive $\mathrm{d}P_k/\mathrm{d}\cos(\theta_k)$, we use a technique from image analysis called a Hough transform. Hough transforms are used to identify particular structures in images, such as circles or lines. The technique is to scan through parameter space for the shape in question, allow each pixel in the image to "vote" whether it is consistent with a shape with those parameters, and then identify places where a lot of pixels voted yes. As an example, one could parametrize a circle by its center and radius. To identify circles in an image, one would step through every possible set of x-coordinate, y-coordinate, and radius. At each step, all pixels that are consistent with that circle vote "1" and all others vote "0". If a circle exists in the image, many pixels will vote "1" at the appropriate coordinates. By finding the step with the most votes, one finds the best circle in the image (and the number of votes gives a measure of how distinct the circle is). In our case, we are indeed looking for rings of photons detected on the PMT sphere, which from the point of view of the event vertex are circles centered on some angular coordinates. Because we know the physics behind the Cherenkov cone, we even have the advantage of knowing the expected radius of the circle ahead of time (the opening angle of the Cherenkov cone). We also make the simple improvement of using a Gaussian kernel instead of a binary vote in our Hough transform since we expect a blurry Cherenkov ring.

Using a Hough transform has some advantages over other techniques to fit a shape to data. Primarily, it looks for the presence of certain features with no penalty for data points that disagree with the expected shape (as would happen in a least-squares fit). It is also capable of finding multiple instances of the shape in a natural way, and characterizes how prominent each instance is. Some drawbacks of the Hough transform are the quantized parameter steps and the considerable volume of parameter space with multiple parameters. When using a kernel, the width and shape of the kernel can also be somewhat arbitrary.

Our procedure to extract $\mathrm{d}P_k/\mathrm{d}(\cos\theta_k)$ from the MC data was as follows. First, we transformed the data into coordinates relative to the true vertex and only examined photons that arrive within 5 ns of the expected arrival time (to remove obviously reflected or scattered photons). We then calculated the Hough transform of the hit coordinates and found the maximum in the voting function, which we considered to be the "best" cone. We then histogrammed the distribution of points as a function of $\cos\theta_k$, the angle relative to the axis of the cone. Repeating this process for many events, we filled a histogram and then fit it with an analytic function.

As expected, the size of the Gaussian kernel affects the shape of $\mathrm{d}P_k/\mathrm{d}(\cos\theta_k)$ that we derive. However, this dependence can be characterized as an underlying shape plus whatever Gaussian we used as a kernel. The underlying shape consists of a component flat in $\cos\theta_k$, plus a broader bump in the cone direction, as shown in Figure 3.3. We fit the distribution with this formula:

$$\frac{\mathrm{d}P_k}{\mathrm{d}(\cos\theta_k)} = (1-f_f)(1-f_p)\,\mathrm{gaus}(\cos\theta_k|\cos\theta_c, w) + (1-f_f)f_p\,\mathrm{para}(\cos\theta_k|0.4, 1.2) + \tfrac{1}{2}f_f \quad (3.5)$$

where gaus$(x|a,b)$ is a Gaussian[2] centered on a of width b, para$(x|a,b)$ is a normalized parabola with roots a and b (and that is 0 in the unphysical region) [3]. f_f is the probability that a Cherenkov photon will be in the flat part of the distribution, and f_p is the probability a photon not in the flat part will be in the parabolic part. We do not worry too much about there being mismatch between the parabolic part and the raw distribution because this light

---

[2] $\mathrm{gaus}\,(x|a,b) = \frac{1}{\sqrt{2\pi}b}e^{-(x-a)^2/(2b^2)}$

[3] $\mathrm{para}\,(x|a,b) = \begin{cases} 0 & x < a \text{ or } x < -1 \\ -A(x-a)(x-b) & 0, a < x < b, 1 \\ 0 & x > b \text{ or } x > 1 \end{cases}$ where $A = 13.8889$ for $a = 0.4$, $b = 1.2$, the only values of a and b we use.

is poorly associated with a cone anyway, so we would prefer the PMT hits to fall within the Gaussian part. It is convenient to parametrize the fractions this way because we can pull out the flat part such that Equation 3.5 becomes

$$\begin{aligned}\frac{\mathrm{d}P_{\text{Cherenkov}}}{\mathrm{d}\Omega}(\theta,\phi) = \frac{1}{4\pi}f_f + \\ (1-f_f)\frac{1}{2\pi}\sum_k f_k\left[(1-f_p)\,\mathrm{gaus}(\cos\theta_k|\cos\theta_c,w) + f_p\,\mathrm{para}(\cos\theta_k|0.4,1.2)\right].\end{aligned} \quad (3.6)$$

Separating the flat part this way is convenient because the flat portion does not depend on the cone direction, so it would be redundant to have each cone contain this part. It also helps the fitter find the correct vertex quickly because the flat part can be large during the burn-in, even though we expect it to be small for the correct fit. Therefore, we allow the flat fraction to float with a constraint that we expect it to be close to zero (we use[4] $\exp(f_f|0.05)$). This constraint ensures that as much light as possible fits to the cones rather than the flat part of the distribution

The parabola in Figure 3.3 is approximately how we would expect additional Cherenkov cones to appear on a plot centered on the best cone (after all, the extra cones do not share their axis with the "best" cone, so photons in these cones will just look like a big blur). We were unable to directly extract separate distributions for multiple cones in the same event because the Hough transform does not provide any indication which photons are accounted for by each cone. However, this is a more reasonable model of a single cone than Figure 3.2. Also, this simple parameterization of $\mathrm{d}P_k/\mathrm{d}\cos\theta_k$ allows us to control the photon distribution while allowing parts of it to float in the MCMC, such as the fraction of light in the flat portion and the fraction in each cone.

Now that we have produced an expected spatial distribution for individual Cherenkov cones, we can use this to fit a MC event with multiple cones and ostensibly unknown vertex position. By fitting for multiple cones, we are able to produce somewhat better fits with our analysis (as measured by recovering the true Monte Carlo parameters more accurately than

---

[4] $\exp(x|a) = \frac{1}{a}e^{-x/a}$

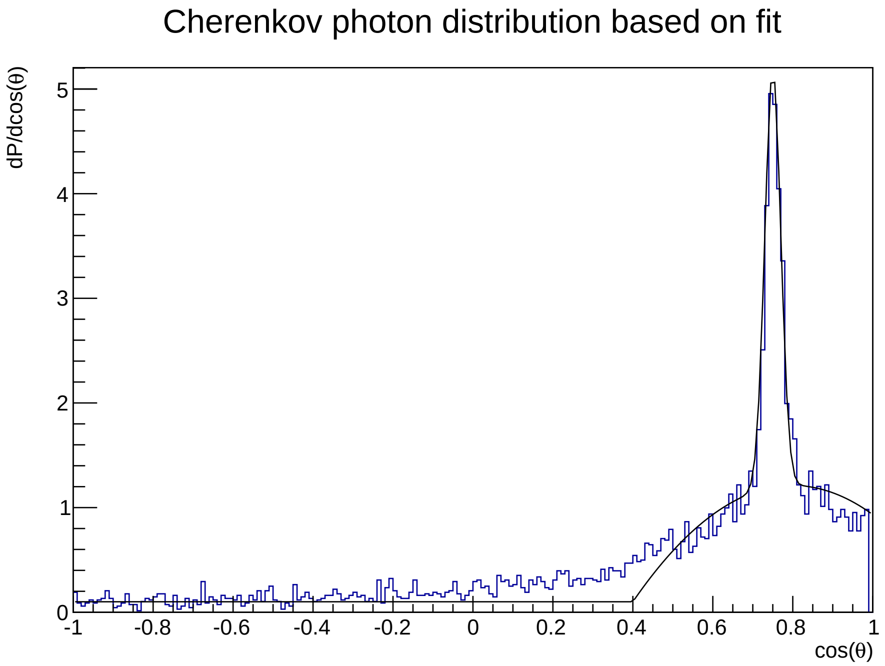

Figure 3.3: Distribution of Cherenkov photons from an electron in water, relative to the electron's fit intial direction, including a fit.

other SNO fitters, see Chapter 4). Each cone has the same vertex $\vec{r}$ (since we expect only a few centimeters of total path length for the electron), but is free to point in any direction. We label the cone axis directions $\vec{u_k}$. We have no prior constraints on $\vec{r}$ or $\vec{u_k}$.

As indicated by Equation 3.6, each cone has a sharp Gaussian and broad parabolic part. We fix the parabolic fraction at $f_p = 0.07$. The width of the narrow Gaussian part of the cone is fixed at $w = 0.10$. The fitter works well for other values of these parameters, or indeed if they are allowed to float, but these values are sufficiently small to ensure it finds well-defined, narrowly-peaked cones.

Time distribution

The temporal distribution for Cherenkov photons is assumed to be a Gaussian of the form

$$\frac{\mathrm{d}P_{\text{Cherenkov}}}{\mathrm{d}t_{\text{PMT}}} = \text{gaus}\left(t_{\text{PMT}}|t_{\text{expected}}, \sigma_t\right) \tag{3.7}$$

where t_{PMT} is the time the photon is detected by the PMT and t_{expected} is the expected arrival time for a prompt photon and σ_t is the spread in time for prompt photons. The expected time is calculated from

$$t_{\text{expected}} = t_{\text{event}} + t_{\text{travel}} \tag{3.8}$$

where t_{event} is the start time of the event (a free parameter) and t_{travel} is the photon's travel time (calculated as $|\vec{q}|/c$, where $\vec{q}$ is the vector from the vertex of the Cherenkov cones to the PMT and c is the average speed of light in the detector, 21.8 cm/s). There are multiple causes for the spread in arrival times σ_t. Photons striking different parts of the PMT will produce pulses at slightly different times. Photons detected after reflecting off of the concentrators (see Chapter 2) will have had a slightly longer travel time. There could also be small amounts of error introduced by the electronics because each channel is not exactly identical. Finally, details like lensing from the AV and finite electron path length, which we normally neglect, will result in a variable spread in the arrival times of the photons, so we allow σ_t to float in each fit. However, by using SNOMAN we have a fairly good idea that σ_t should be 1.7 ± 0.3 ns, so we include this as a prior.

3.3.2 Reflected light

The second component of our model is photons reflected or scattered before being detected. These photons will all arrive later than the prompt photons, with a piecewise $\mathrm{d}P_{\text{reflected}}/\mathrm{d}t_{\text{PMT}}$ shown in Figure 3.4 that was constructed based on the output of the SNOMAN Monte Carlo. SNOMAN, in turn, was tuned from calibration data using a source that produced reasonably isotropic light in short bursts, so that PMT-to-PMT variations could be characterized. The fact $\mathrm{d}P_{\text{reflected}}/\mathrm{d}t_{\text{PMT}}$ has so much structure (e.g. a minimum at 15 ns and a peak at 36 ns)

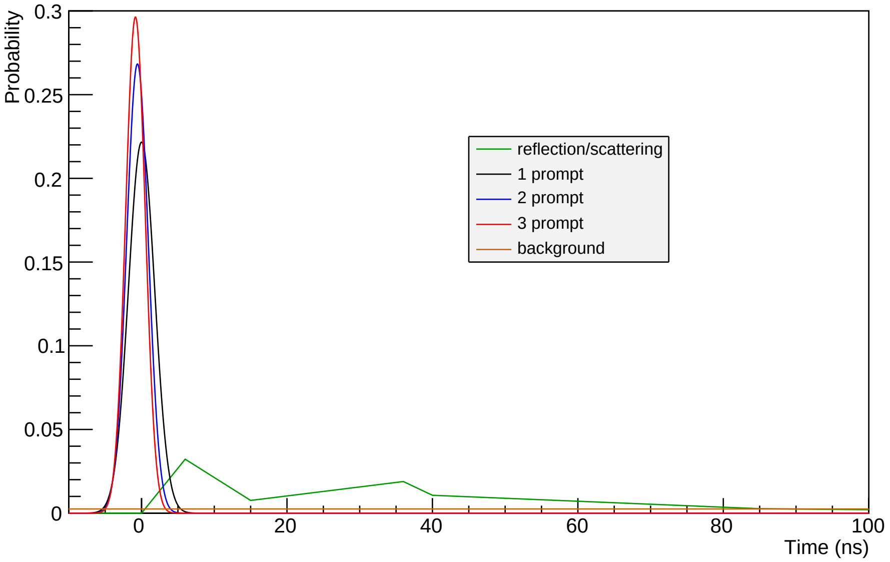

Figure 3.4: The temporal distributions $\mathrm{d}P/\mathrm{d}t(t|n_j)$ used in the fitter. When multiple Cherenkov photons reach the same PMT, the earliest determines the detection time.

is due to reflections off the AV or PMTs having to travel reasonably well-defined distances before being detected on the other side of the detector.

$$\frac{\mathrm{d}P_{\text{reflected}}}{\mathrm{d}t_{\text{PMT}}} = \begin{cases} 0 & t_{\text{PMT}} \leq 0 \\ 0.0053714 t_{\text{PMT}} & 0 < t_{\text{PMT}} \leq 6 \\ -0.02737 t_{\text{PMT}} + 0.04865 & 6 < t_{\text{PMT}} \leq 15 \\ 0.00053763 t_{\text{PMT}} - 0.00046925 & 15 < t_{\text{PMT}} \leq 36 \\ -0.0020528 t_{\text{PMT}} + 0.0927858 & 36 < t_{\text{PMT}} \leq 40 \\ -0.00017791 t_{\text{PMT}} + 0.0177904 & 40 < t_{\text{PMT}} \leq 85 \\ -0.000041055 t_{\text{PMT}} + 0.006158575 & 85 < t_{\text{PMT}} \leq 150 \\ 0 & 150 < t_{\text{PMT}} \end{cases} \tag{3.9}$$

We assume the spatial distribution of the reflected light is isotropic with respect to the center of the detector. The reason for this is that despite the spherical symmetry of the detector, the math is actually pretty complicated for where light reflecting off the AV will end up.

Even worse, the PMTs and concentrators have extremely complicated shapes, so reflections from these surfaces can go basically anywhere. Since the reflections are complicated, and reflected photons are already easily identified by their arrival times, the spatial information would give very little information on the event parameters anyway.

3.3.3 Backgrounds

The simplest part of the physical model is the backgrounds. This is assumed to be PMT noise, independent of any parameters of the physics event. The backgrounds are evenly distributed in space (4π steradians) and time (the 400 ns of the event):

$$\frac{\mathrm{d}P_{\text{background}}}{\mathrm{d}\Omega}(\theta, \phi) = \frac{1}{(4\pi)(400\ \text{ns})} \tag{3.10}$$

3.4 Expected number of photons at a PMT

In order to fit an event, we will calculate the expected number of photons at PMT i from distribution j, λ_{ij}, where j is over Cherenkov, reflection/scattering, and background.

The simplest case is for $j =$ background, which is the same at each PMT:

$$\lambda_{i(j=\text{background})} = \frac{b}{N_{\text{PMTs}}} \tag{3.11}$$

where b is the number of background hits in the whole detector for this event.

For the prompt and reflected light distributions we can write

$$\lambda_{ij} = N f_j \frac{\mathrm{d}P_j}{\mathrm{d}\Omega}(\theta_i, \phi_i|\Psi) P_{ij}(\text{detection}|\Psi) \tag{3.12}$$

where N is the total number of photons we expect in this event (a measure of the energy of the event), f_j is the fraction of the photons in each distribution, $\mathrm{d}P_j/\mathrm{d}\Omega(\theta_i, \phi_i|\Psi)$ is the probability for a photon of type j to travel in the direction of PMT i, and $P_{ij}(\text{detection}|\Psi)$ is the probability of such a photon actually reaching the PMT and being detected.

We already calculated explicit forms for the factor $\mathrm{d}P_j/\mathrm{d}\Omega(\theta_i, \phi_i|\Psi)$ in the previous section. For example, we treat the the reflected light as isotropic (from the center of the detector), and the Cherenkov light follows a distribution given by Equation 3.6, relative to the vertex of the cones.

The factor $P_{ij}(\text{detection}|\Psi)$ includes several things. Most important is the solid angle subtended by the PMT. For the prompt light, this is the PMT's surface area divided by the distance between the event position and PMT, squared, times the foreshortening of the PMT as viewed from the event position. Since a full sphere is 4π steradians, to create a probability we divide the solid angle by 4π. This part of $P_{ij}(\text{detection}|\Psi)$ is then

$$\frac{A_{\text{PMT}}}{4\pi\vec{q}^2}(\hat{q}\cdot\hat{x}_i) \tag{3.13}$$

where $\vec{q}$ is the vector from where the light was generated to the PMT and $\vec{x_i}$ is the vector from the center of the detector to the PMT (so that $\hat{x}_i$ is an outward-pointing unit vector). For reflected light, we assume an isotropic distribution, so we just assume light was created at the center of the detector by substituting $\vec{q} = \vec{x_i}$.

$P_{ij}(\text{detection}|\Psi)$ also includes the PMT's (and concentrator's) angular acceptance. This is approximated (based on measurements) to be a step function, with photons approaching at greater than 0.96 radians from normal being completely attenuated.

Finally, we also consider the probability a photon is absorbed en route. We individually consider the absorption length in D_2O and H_2O. We treat the absorption in the acrylic vessel as an adjustment to the absorption in the H_2O, since the path length through the acrylic is highly correlated with that through the H_2O. It turns out the absorption length for light in water is highly dependent on the purity of the water, so it differs pretty strongly between the two different water regions and varies over time. SNO made measurements of absorption lengths in situ throughout the experiment, and from these measurements we decided to use lengths of 100 m in H_2O and 50 m in D_2O.

For the Cherenkov photons, calculating the path lengths in the H_2O and D_2O is a straightforward geometry problem. For the reflected light, we imagine a photon starting somewhere random in the detector, reflecting off a PMT or the AV, and then being detected by a PMT on the far side of the detector. The average chord length of a sphere is 4/3 its radius, so if a photon travels this length on average one and a half times, it passes through 1200 cm of D_2O. Meanwhile, if the photon reflects off a PMT, it passes through the H_2O three times, but if it reflects off the AV, it only passes through the H_2O once. On average then, a reflected photon travels twice through the D_2O, which is a total of about 600 cm.

All together, then,

$$\begin{aligned} P_{i(j=\text{Cher})}(\text{detection}|\Psi) &= \left[\frac{A_{\text{PMT}}}{4\pi\vec{q}^2}(\hat{q}\cdot\hat{x}_i)\right]\left[\Theta(0.96-\cos^{-1}(\hat{q}\cdot\hat{x}_i))\right]\left[e^{-q_D/l_{\text{abs,D}}-q_H/l_{\text{abs,H}}}\right] \\ P_{i(j=\text{reflected})}(\text{detection}|\Psi) &= \frac{A_{\text{PMT}}}{4\pi\vec{x}_i^{\,2}}e^{-1200/l_{\text{abs,D}}-600/l_{\text{abs,H}}}. \end{aligned} \tag{3.14}$$

3.5 *The likelihood function*

We now have all the tools we need to derive the likelihood function $P(\Xi|\Psi)$. Each PMT is independent of the others, so the overall event likelihood is the product of the likelihoods for each PMT:

$$P(\Xi|\Psi) = \prod_{i=\text{all PMTs}} \frac{\mathrm{d}P}{\mathrm{d}t_i}(n_i, t_i|\Psi), \tag{3.15}$$

where n_i is the number of photons collected by PMT i, and t_i is the time at which the hit was registered.

It turns out that PMTs suffer from a technical limitation called “dead time”. When multiple photons arrive at a PMT within about 30 to 40 ns, only one pulse is detected. (The charge registered by the PMT does increase, giving a hint of the number of photons, but this is not reliable enough to be sure how many.) Since the duration of an event is only a few nanoseconds, this means that we cannot know n_i ahead of time, except to say that there were no hits, or there was at least one hit. In most cases, we expect multiple photons incident on one PMT to be rare, but we do have to make some considerations for the possibility that a PMT hit may represent multiple photons, otherwise, we may accidentally penalize a PMT for “only” detecting one photon. It also means that we must consider t_i to be the time at which the *first* photon was detected by that PMT.

Recall that in our model, photons may come from three possible sources: Cherenkov radiation, reflections (or scatters), or background. Therefore, we consider n_i to be in fact $\vec{n_i}$, a vector containing n_{ij}, the number of photons at PMT i from distribution j. Because we don't actually know n_{ij}, we will sum over all possibilities of $\vec{n_i}$ with $n_i > 0$ for a PMT that registered a hit (or use $\vec{n_i} = \vec{0}$ if the PMT registered no hit).

Now we wish to calculate $\mathrm{d}P/\mathrm{d}t_i(n_i, t_i|\Psi)$. First we break our function into a sum of all possible values of $\vec{n_i}$ consistent with the data.

$$\frac{\mathrm{d}P}{\mathrm{d}t_i}(n_i, t_i|\Psi) = \sum_{\vec{n_i}} \frac{\mathrm{d}P}{\mathrm{d}t_i}(\vec{n_i}, t_i|\Psi) \tag{3.16}$$

Then, we expand the two-variable joint probability into a conditional probability.

$$\frac{\mathrm{d}P}{\mathrm{d}t_i}(\vec{n_i}, t_i|\Psi) = \frac{\mathrm{d}P}{\mathrm{d}t_i}(t_i|\vec{n_i}, \Psi)P(\vec{n_i}|\Psi) \tag{3.17}$$

Now, we can be more explicit using the fact that the probability of detecting a photon at a certain time depends on the numbers and types of photons present, n_{ij}, t_{event}, and σ_t, but not on any other parameters.

$$\frac{\mathrm{d}P}{\mathrm{d}t_i}(t_i|\vec{n_i}, \Psi) = \frac{\mathrm{d}P}{\mathrm{d}t_i}(t_i|\vec{n_i}, t_{\text{event}}, \sigma_t) \tag{3.18}$$

The factor $P(\vec{n_i}|\Psi)$ in Equation 3.17 is given directly by Poisson statistics. We simply calculate $\lambda_{ij}(\Psi)$ as described in the previous section, and then:

$$P(\vec{n_i}|\Psi) = P(\vec{n_i}|\lambda_{ij}) = \prod_j \frac{e^{-\lambda_{ij}}(\lambda_{ij})^{n_{ij}}}{n_{ij}!} \tag{3.19}$$

The factor $\mathrm{d}P/\mathrm{d}t_i(t_i|\vec{n_i}, t_{\text{event}}, \sigma_t)$ in Equation 3.18 is the temporal distribution of the first photon detected given a known number of photons from each distribution.

We calculate these on a case-by-case basis in Table 3.1 (N.B. several of these are shown in Figure 3.4).

In the case of no photons, $\mathrm{d}P/\mathrm{d}t_i(t_i|\vec{n_i} = \vec{0}, t_{\text{event}}, \sigma_t)$ is somewhat ill-defined, since we never detect a photon. We use 1 so that the likelihood reduces to the likelihood of simply detecting no photons at that PMT.

$n_{i(j=\text{background})}$	$n_{i(j=\text{prompt})}$	$n_{i(j=\text{reflected})}$	$\mathrm{d}P/\mathrm{d}t_i(t_i\vert\vec{n_i})$
0	0	0	1 (no photon detected)
0	0	1	$\mathrm{d}P/\mathrm{d}t(t\vert n_{j=\text{refl}}=1)$ (Fig 3.4)
0	1	any	$\mathrm{d}P/\mathrm{d}t(t\vert n_{j=\text{prompt}}=1)$ (a Gaussian)
0	2	any	different Gaussian (see Fig 3.4)
0	3+	any	different Gaussian (see Fig 3.4)
1	0	0	1/400 (uniform over event window)
1	0	1	uniform if $t<0$, then $\mathrm{d}P/\mathrm{d}t(t\vert n_{j=\text{refl}}=1)$
1	1	any	uniform if $t<0$, then $\mathrm{d}P/\mathrm{d}t(t\vert n_{j=\text{prompt}}=1)$

Table 3.1: Temporal PDFs with each considered combination of photon sources.

In the case where there is only one photon, $\mathrm{d}P/\mathrm{d}t_i(t_i|\vec{n_i}, t_{\text{event}}, \sigma_t)$ is simply $\mathrm{d}P_j/\mathrm{d}t$ as derived in Section 3.3.

Values of $\vec{n_i}$ with multiple photons require more consideration. First of all, the expected rate of background and reflected photons at any PMT is so low that we don't have to seriously consider $n_{ij} > 1$ cases (we simply replace $P(n_{ij} = 1|\lambda_{ij})$ with $P(n_{ij} > 0|\lambda_{ij})$ for $j \in \{\text{background, reflected}\}$ and assume the temporal distribution is the same as in the case of $n_{ij} = 1$). On the other hand, the Cherenkov light is sometimes quite concentrated, such that some PMTs have $\lambda_{i(j=\text{Cherenkov})}$ as high as 1, which results in a realistic probability of as many as 3 photons incident on those PMTs. In these cases, we expect the first photon to be somewhat earlier than we would expect with just one photon, as shown in Figure 3.4. The exact distribution is difficult to calculate analytically, but fortunately it is simple to find numerically as the earliest photon when multiple are drawn from the same distribution. The result fits well to a different Gaussian for each value of $n_{i(j=\text{Cherenkov})}$.

In cases with both prompt and reflected photons reaching a PMT, we can simply use the prompt distribution since this photon will always arrive earlier than a reflected one.

A background "photon" could, rarely, appear early and mask a prompt photon. A fair approximation for this distribution is to use $\mathrm{d}P_{j=\mathrm{background}}/\mathrm{d}t_i$ while $t_{\mathrm{PMT}} < t_{\mathrm{expected}}$, then $\mathrm{d}P_{j=\mathrm{Cherenkov}}/\mathrm{d}t_i$ (scaled to the correct normalization) while $t_{\mathrm{PMT}} > t_{\mathrm{expected}}$. The same trick works with one background and one reflected photon.

3.6 List of parameters and priors

At this point it will be useful to summarize all the parameters and priors in the event fitter. Recall that a single PMT's likelihood depends on λ_{ij}, which in turn depend on the event position $\vec{r}$, the cone directions $\vec{u_k}$, the number of background hits b, the number of event hits N, and the fractions f_f, f_k, and $f_{j=reflected}$. The PMT's likelihood also depends on the time of the event t_{event} and the width of the prompt time distribution σ_t.

Additionally, a number of parameters are fixed. We fix the speed of light c, at 21.8 cm/s, and use the same value for both heavy and light water. We also fix the parameters w, θ_c, and f_p from the Cherenkov distribution, and the absorption lengths in H_2O and D_2O. Table 3.2 summarizes all of the floating parameters and the priors we place on them.

parameter	description	enters fitter in	prior	step sizes
N	expected number of event photons	$\lambda_{i(j=\text{Cherenkov,reflected})}$	uniform	1 photon
t_{event}	time of the event	$\mathrm{d}P_{j=\text{Cherenkov,reflected}}/\mathrm{d}t_i$	uniform from 0-400 ns	0.2 ns
$\vec{r}$	event position	$\lambda_{i(j=\text{Cherenkov})}$	uniform in x, y, and z	1.5 cm
$\vec{u_k}$	direction of cone k	$\lambda_{i(j=\text{Cherenkov})}$	uniform in all directions	0.05 rad
$f_{j=\text{reflected}}$	fraction of light reflected	$\lambda_{i(j=\text{reflected})}$	uniform from 0-1	0.005
f_f	fraction of non-cone Cherenkov light	$\lambda_{i(j=\text{Cherenkov})}$	$\exp\left(f_f\|0.05\right)$	0.02
f_k	fraction of Cherenkov light in cone k	$\lambda_{i(j=\text{Cherenkov})}$	uniform from 0-1	0.05
σ_t	width of $\mathrm{d}P_{j=\text{Cherenkov}}/\mathrm{d}t_i$	$\mathrm{d}P_{j=\text{Cherenkov}}/\mathrm{d}t_i$	$\text{gaus}\left(\sigma_t\|1.7, 0.3\right)$	0.045 ns
b	background PMT rate	$\lambda_{i(j=\text{background})}$	$\text{gaus}\left(b\|3, 1\right)$	0.2 hits

Table 3.2: Fitter parameters, descriptions, and priors

3.7 Background discrimination

Apart from finding the energy, position, and other parameters of an event, we also wish to determine the probability that an event was of a certain *type*. In the previous sections of this chapter, we discussed fitting events that were produced by relativistic electons that produce Cherenkov radiation, such as caused *hep* and ^{8}B neutrinos. Other sources of relativistic electrons are Michel electrons (products of muon decay), elastic scattering of atmospheric neutrinos, and Compton scatter of γs.

However, not all events are of this type, and using the fitter we can find the probability that an event is of a different type.

What we wish to find is $P(T|\Xi)$, the probability that the event was type T given the data. Bayes's Theorem tells us

$$P(T|\Xi) = \frac{P(\Xi|T)P(T)}{P(\Xi)}. \tag{3.20}$$

Alone, this is not very useful because we do not know $P(T)$ (the probability an event was a certain type) or $P(\Xi)$ (the probability of the observed data). However, let us take the ratio between the probabilities that an event was one of two types T and T':

$$\frac{P(T|\Xi)}{P(T'|\Xi)} = \frac{P(\Xi|T)}{P(\Xi|T')}\frac{P(T)}{P(T')} \tag{3.21}$$

Now $P(\Xi)$ has canceled out. We still have the ratio of the prior probabilities of each type, $\frac{P(T)}{P(T')}$ (which we do not know), but this is a constant that does not depend on the data. Equation 3.21 then tells us that the ratio of the probabilities of an event being of two types is proportional to the ratio of the likelihoods of those types.

The strategy to use Equation 3.21 then is to separately fit the event assuming it was of each type, then take the ratio of the likelihoods calculated for each type. We can choose a cutoff value for this ratio, and events with a higher ratio are classified as type T and events with a lower ratio are classified as type T'. Choosing this cutoff value is equivalent to choosing a prior for $\frac{P(T)}{P(T')}$, though it also corrects for the fact that event types using more parameters would otherwise unfairly improve the likelihoods for those types.

The fits for each event type are done in exactly the same way as described earlier in this chapter, only with different distributions for $\mathrm{d}P_j/\mathrm{d}\Omega(\theta_i, \phi_i|\Psi)$. The "likelihood" we use, as mentioned in Section 3.2, is actually the likelihood times the priors, which is calculated at each step in the MCMC. For comparing event types, we use the mean of the posterior of the logarithm of the likelihood ratio.

3.7.1 Acrylic vessel instrumental backgrounds

To be specific, the type of event we wish to identify with this test is called "acrylic vessel instrumental backgrounds" or AVIBs[5]. AVIBs are a class of events of a somewhat mysterious origin that were problematic for SNO[6]. In previous analyses, AVIBs reconstruct to the position of the AV and occur over a range of energies. Monte Carlo simulations of these events have been done, and the simulations that produce events best matching the data are those that generate photons isotropically from points near the inner boundary of the AV. These Monte Carlo events match the real events well whether they are generated just within the bulk acrylic or just within the D_2O. These photons are generated with wavelengths of 335-500 nm [38]. This wavelength range covers basically all the possibilities because shorter wavelengths are completely absorbed by the acrylic, and longer wavelengths are not detected by the PMTs.

The spatial distribution $\mathrm{d}P/\mathrm{d}\Omega(\theta, \phi|\Psi)$ we use for AVIBs is

$$\frac{\mathrm{d}P_{\mathrm{AVIB}}}{\mathrm{d}\Omega}(\theta_r, \phi_r) = \frac{1}{2\pi} 0.764602 \sqrt{|\cos(\theta_r) + .07)|} \tag{3.22}$$

where θ_r is the angle at which the photon leaves the acrylic, relative to the radial direction of the event. This distribution is shown in Figure 3.5. This form is fit from a simple model

---

[5] Within the SNO collaboration these events are informally called "Leslie events" after their discoverer.

[6] My pet theory is that these events are triboluminescence in the acrylic. A few hundred small cracks have been observed in the inner surface of the AV, localized in a way consistent with the AVIBs. However, there was not a detailed survey of the crack locations so I cannot compare the locations of the cracks and AVIBs. Furthermore, even if this were confirmed as the source of the events, it wouldn't be that informative since measurements of triboluminescence of acrylic submerged in heavy water do not appear to have been conducted. For what it is worth, the spectrum of triboluminescence typically resembles an electric discharge.

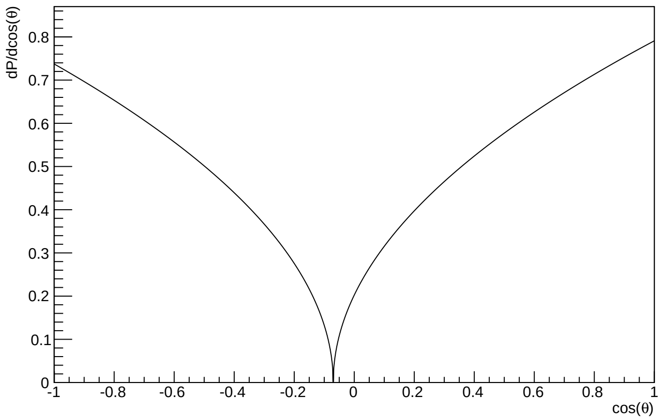

Figure 3.5: Spatial distribution of photons from AVIB events.

that just traces rays (with refraction and reflection) from random points within a spherical AV until they reach an 824-cm radius, which fits well to a square-root function. The shift of 0.07 is due to the finite thickness of the AV.

In an AVIB event, as in a Cherenkov event, some of the light will still scatter or reflect before being detected. In fact, due to the spherical geometry, if the event actually does occur inside the acrylic we may expect it to undergo total internal reflection several times, though imperfections in the spherical shape of the acrylic will allow it to escape eventually. Therefore, we include the same reflection distribution as in the Cherenkov events. Also, the "background" PMT dark rate is expected to be the same regardless of the event type.

The temporal distribution for the AVIBs is fit from the SNOMAN Monte Carlo data for AVIBs. The functional form is

$$\frac{\mathrm{d}P_{\mathrm{AVIB}}}{\mathrm{d}t_i} = 0.0996724 \times \begin{cases} e^{-(t_i^2)/2(2.42725)^2} & t_i < 2.95655\ ns \\ e^{-0.143011 t_i - 0.345818} & 2.95655\ ns < t_i < 14.4765\ ns \\ e^{-0.0457376 t_i - 1.72445} & t_i > 14.4765\ ns \end{cases} \tag{3.23}$$

where t_i is the expected arrival time. When we run the fits, we include the above AVIB distribution, the same reflected (or scattered) light distribution as for Cherenkov photons,

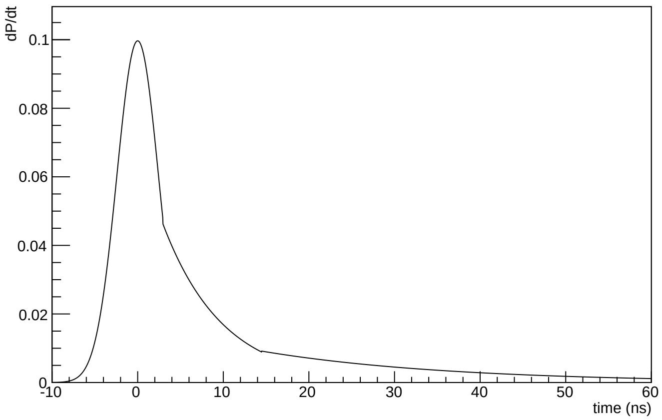

Figure 3.6: Probability distribution for photons from AVIB events.

and the same background PMT rate. The distribution is shown in Figure 3.6.

3.7.2 Cherenkov events for comparison

Most of this chapter was dedicated to Cherenkov events, so I will not dwell on them here. However, it does bear mention that while the distribution derived earlier fits Cherenkov events very well, it also fits non-Cherenkov events a little too well. For example, large values of f_f allow isotropic events to be fit, and two cones pointed in opposing directions do a pretty good job of fitting an AVIB event.

For the purpose of identifying Cherenkov-like events while rejecting AVIBs, we ran the Cherenkov fitter a second time with only one cone. This gives a good fit for electron-like events but a bad fit for AVIBs, as desired.

Chapter 4

EVENT FITTER PERFORMANCE

In this chapter, the performance of the fitter described in Chapter 3 (which we will refer to as "Cone Fitter") will be evaluated. This evaluation will be conducted both relative to existing fitters, and on an absolute basis where possible. We will make particular use of SNOMAN Monte Carlo, which is detailed in Chapter 2.

Cone Fitter keeps track of many possibly useful variables at each step, including the position of light production, number of photons produced, the direction and relative intensity of each cone, the fraction of the distribution that is late or not associated with cones, and the width of the prompt time distribution. While each of these pieces of information may give useful information about an event, the most interesting and physically relevant outputs are the position and energy (which in our case is the number of photons) of the event. A subtlety is that Cone Fitter actually returns a *posterior distribution* for each variable, rather than a single value. For this chapter, we will use the mean of the posterior distribution as the best value.

4.1 Previous fitters

Before comparing Cone Fitter to other fitters, let us first discuss those other fitters. It turns out Cone Fitter is somewhat unique among SNO fitters in that it fits for position and energy simultaneously. In previous SNO analyses, one fitter would first find the position of the event which would then be used as input for a second fitter to find the energy of the event.

Several different position fitters were used in SNO. Most were quite simple, for instance minimizing the time residuals between expected and actual detection of the prompt light. Some were more complicated, and arguably the best position fitter, called "Path Fitter" or “FTP” improved on previous fitters by considering possible actual paths between the vertex and each hit PMT, rather than making assumptions to simplify this calculation. Few fitters made use of the spatial distribution of events, and those that did often suffered from a bias attributed to the diffuse nature of an electron-based Cherenkov cone. FTP was very successful in part because it did use the shape and direction of the Cherenkov cone to improve the fit position. For details on Path Fitter, see [27].

SNO energy fitters largely worked by counting the number of prompt photons, and often making some adjustments based on various considerations like the position of the event or the geometry of the detector. SNO’s second-best energy fitter, RSP, worked on this basis, using only prompt photons and information about the detector geometry and optics to calculate the event’s energy. The most recent iteration of RSP also takes into account PMT-to-PMT variations to improve the energy estimate [29].

For Phases I and II, the best energy fitter was called FTK. FTK uses the position returned by Path Fitter and runs a mini Monte Carlo for that event in order to find the conditions most like the real event [19]. One big improvement FTK made was using the late light in addition to the prompt light, by taking careful account of reflections and scattering. FTK also carefully integrated over all wavelengths, calculated all path lengths (for scattering and absorption), and used detailed detector geometry including the neck, belly plates, and ropes.

FTK has a few drawbacks, however, which motivated the creation of Cone Fitter. One is that many physics events have multiple cones, for instance if an electron undergoes a large scatter. FTK also goes haywire if Path Fitter returns a position inside the bulk of the Acrylic Vessel because real events in acrylic produce little light, so an event with a typical amount of light that is reconstructed within the acrylic is inferred to have been very energetic. This was not a large drawback previously because other analyses had conservative fiducial volume cuts. Finally, FTK was never adapted to Phase III, when the Neutral Current Detectors (an array of metal tubes) were installed inside the D_2O, complicating the Monte Carlo part of the fit. Therefore, RSP remained the best energy fitter during Phase III.

4.2 Position performance

A simple test for how well each fitter reconstructs the position of an event is to take Monte Carlo data (for which you know the true position of the event), fit each event with each fitter, then plot histograms of the deviation between the fit position and true position. The Monte Carlo events used in these fits were generated as *hep* charged current events, meaning the electrons extend up to a maximum of about 18 MeV, with a distribution peaked around 11 MeV. *hep* events are a sensible choice because, in addition to being the signal, they look exactly like the largest background, ^{8}B events, which all occur within a subset of the *hep* energy range. These events are uniformly distributed in the D_2O and isotropic in direction. Each histogram below will show two distributions for FTP, because FTP produced some figures of merit which were used to cut poorly-fit events (these were a small number of events in this sample). The cuts also included some high-level cuts, including a cut which removes events at a radius higher than 5.5 m, which is responsible for most of the events that are cut. The Cone Fitter distributions look essentially identical with or without these cuts, so no events are cut from the Cone Fitter plots.

We will discuss residuals (that is, difference between the coordinates of the fit position and true position) in a variety of coordinate systems in this section. Because the detector has rotational symmetry about the z-axis, we first consider a cylindrical system with its origin at the center of the AV, a z-coordinate that points upward through the neck of the AV, and $\rho^2 = x^2 + y^2$ is the distance squared from the z-axis. Since the detector is nearly spherically symmetric, we also look for bias in the radial position at which events are reconstructed. We will also show the overall distance between events' true and reconstructed positions. A final coordinate system we consider is oriented with respect to the true direction of the event, $\vec{u}$. The position along the axis defined by $\vec{u}$ is given by $\vec{u} \cdot \vec{r}$, and the distance squared away from this axis is given by $|\vec{u} \times \vec{r}|^2$. Note that the origin of this cylindrical coordinate system is still the center of the detector, just with the axis oriented to the electron's velocity. This system is useful because we may reasonably expect biases in the direction of the initial electron, even if these average out in detector coordinates.

Figures 4.1 through 4.3 show position residuals for Phases I-III, respectively. The residuals for Cone Fitter and FTP in z and the xy-plane are shown in Subfigures (a) and (b) of each figure. Additionally, residuals in the radial direction and overall distance between the fit and true position are shown in Subfigures (c) and (d). Distributions for residuals in the direction parallel to the true electron direction and perpendicular to this direction are also shown in Subfigures (e) and (f).

Let us examine Figures 4.1 through 4.3. In each phase, in detector coordinates, Cone Fitter shows an improved resolution over FTP, of about 1 to 2 cm in each case (roughly a 10% improvement), where resolution is defined as the standard deviation in Subfigures (a), (c), and (e), and the means of Subfigures (b), (d), and (f).

Subfigures (c) in the above plots show that Cone Fitter appears to have a very small radial bias (about 1 cm outward), though this is a comparable size with FTP (1 cm inward). One should note that if errors are isotropic, one expects the average residual in radius to be slightly positive because of purely volumetric considerations (there is more volume near any given point that is farther from the origin than there is nearer to the origin). A simple Monte Carlo estimates that with a 14 cm resolution and a 6 m radius volume, an unbiased fitter should actually have a mean radial residual of +0.5 cm. For the same resolution but with a *fit* radius cut at 550 cm, as applied in the FTP-with-cuts case, the average should be -0.6 cm. Deviations from these values could be caused, for example, by imperfect knowledge of the speed of light in the media.

In Subfigures (e) of the above figures, we can see the residuals are biased in the forward (relative to the electron) direction for each fitter. This is actually expected because the electron travels a finite distance while emitting Chereknov photons, so the best fit position should be roughly the center of this path, a few centimeters forward from the true vertex. The path length varies with energy, and for energies of interest (around 10 MeV), we expect path lengths of around 5 cm, so a forward residual with a mean of 2-3 cm is expected. Additional bias in this value could be caused by imperfect knowledge of the Cherenkov cone

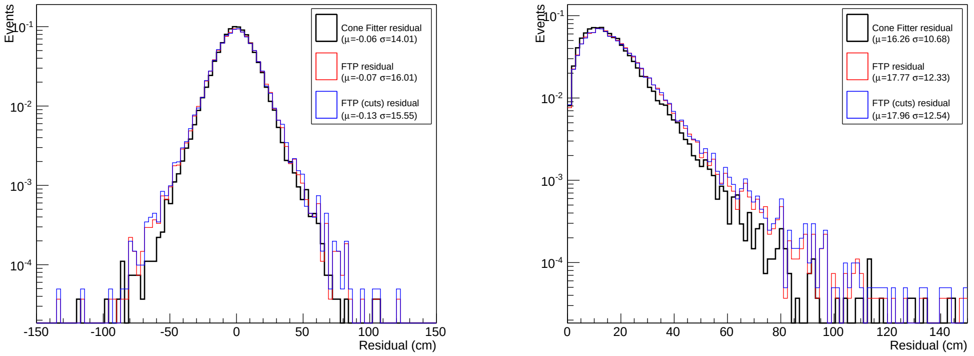

(a) z residuals.

(b) xy residuals.

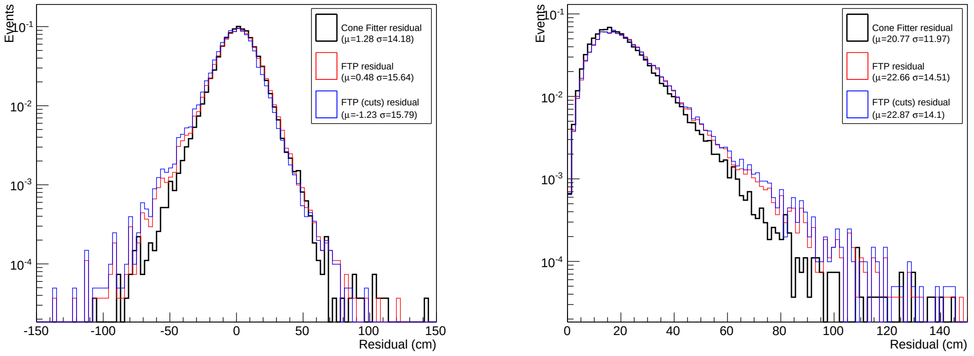

(c) r residuals.

(d) 3D residuals.

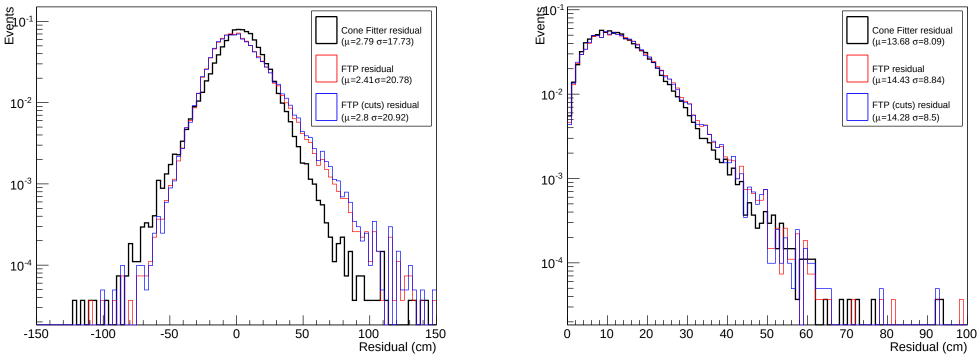

(e) Parallel-to-electron residuals.

(f) Perpendicular-to-electron residuals.

Figure 4.1: Phase I position residual histograms.

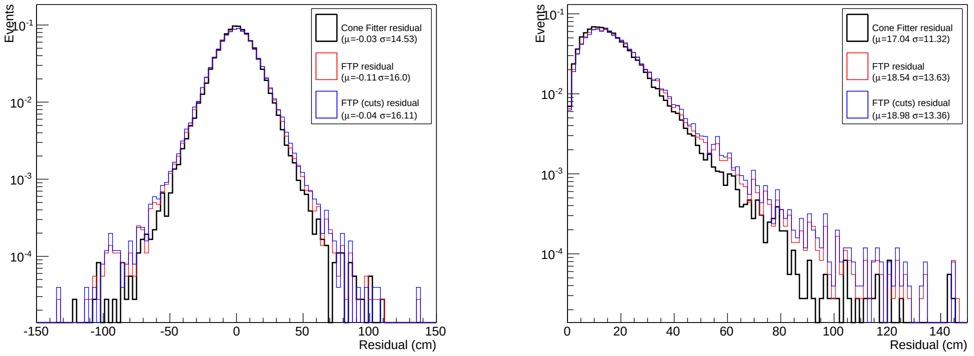

(a) z residuals.

(b) xy residuals.

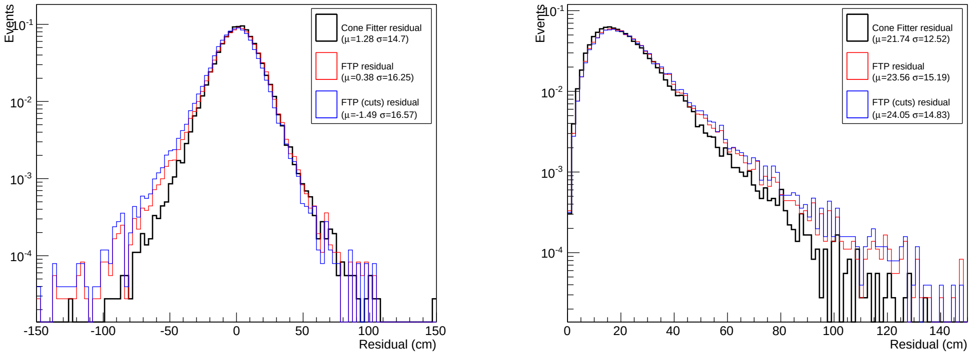

(c) r residuals.

(d) 3D residuals.

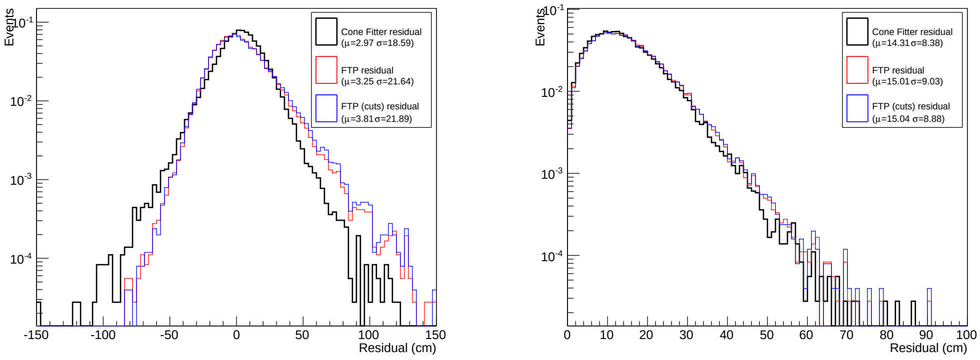

(e) Parallel-to-electron residuals.

(f) Perpendicular-to-electron residuals.

Figure 4.2: Phase II position residual histograms.

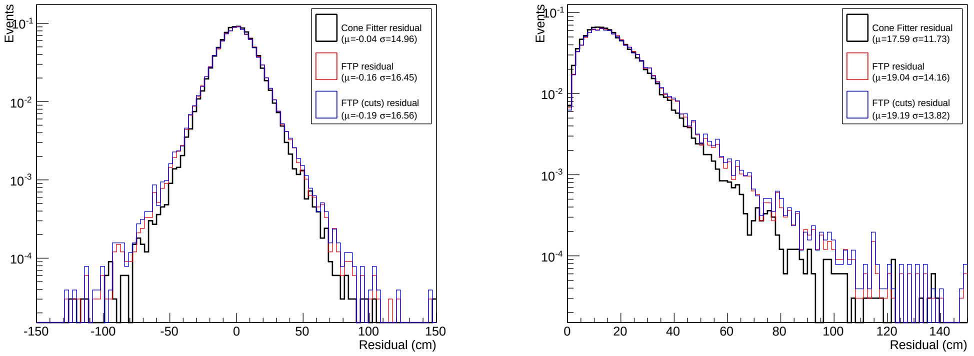

(a) z residuals. (b) xy residuals.

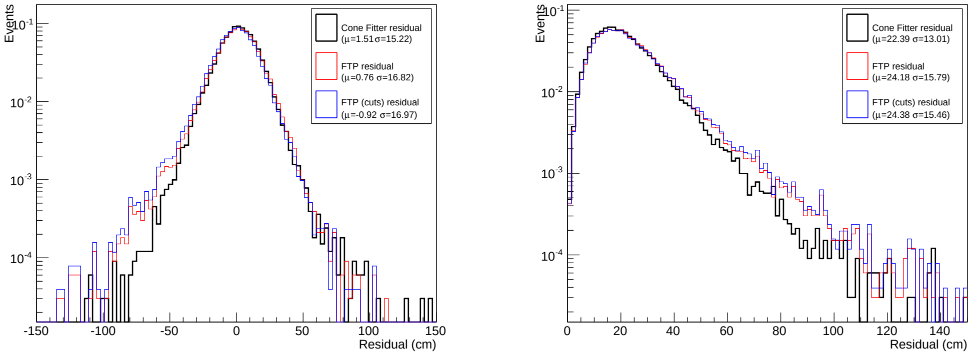

(c) r residuals. (d) 3D residuals.

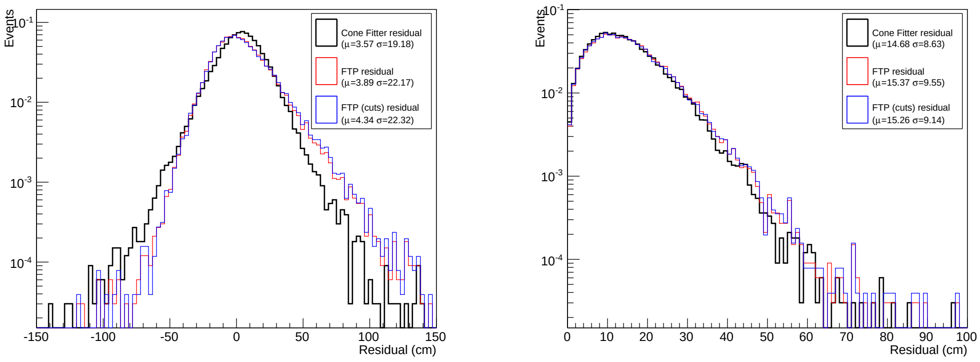

(e) Parallel-to-electron residuals. (f) Perpendicular-to-electron residuals.

Figure 4.3: Phase III position residual histograms.

opening angle, the asymmetrically-diffuse Cherenkov photon distribution, or lensing effects.

Subfigures (e) in the above plots also show a somewhat different behavior between Cone Fitter and FTP. Cone Fitter's residuals tend to be more symmetrical about the center, while FTP has a larger tail in the forward direction than backward. SNO fitters have a tendency to be biased forward along the electron direction (a tendency informally called "drive"). For a discussion of this effect, see for example [27]. Possible reasons Cone Fitter is more symmetric in $\vec{u} \cdot \vec{r}$ are that the full fit (position, direction, and energy) is done simultaneously, Cone Fitter uses a PDF that may not as strongly penalize non-Cherenkov (and in particular, late) light, or perhaps because in this analysis we have used the mean of the posterior distribution, rather than just the maximum likelihood.

4.3 Energy performance

FTK and RSP are both calibrated to return energy values directly, while Cone Fitter has no such internal calibration. Instead, Cone Fitter returns an effective number of photons generated, which must be calibrated to return an energy.

At the end of the day, this calibration will be floated to match the data. This is necessary in case details of the Monte Carlo are biased, such as PMT response, number of Chereknov photons produced, errors in absorption values, etc. We also expect variations in this calibration phase-to-phase, particularly in Phase III when the neutral current detectors were installed in the D_2O region, which will obviously block some photons from being detected. However, even between Phases I and II we may expect some difference in the photons-to-energy relationship as the salt may affect the optical properties of the D_2O. Additionally, over time parts of the detector age in a way that negatively affects its ability to detect photons. A particularly clear example of this is the PMT concentrators (aluminum fins used to reflect additional light into the PMTs), which deteriorated over the course of the experiment.

One example of the detector aging that deserves special mention is that over the course of the experiment, roughly ten percent of the PMTs failed. This should not bias the energy returned by Cone Fitter, because Cone Fitter uses run-by-run data about PMTs with anomalous occupancy (e.g. if a PMT has zero, very few, or far too many hits in the course of a run, compared to other PMTs) and includes every active PMT in the energy calculation, including those that are not hit. Therefore, having fewer working PMTs may increase the uncertainty on an energy, but should not bias it downward.

In order to evaluate the performance of Cone Fitter in calculating energy, the first thing we wish to know is whether the energy parameter is linear, as we expect. We use the same Monte Carlo *hep* event fits as in the position results above. We bin these events by true energy in 1-MeV bins, calculating the mean and standard error of the Cone Fitter energy for each bin. We then fit a line to these points and use the χ^2 value to determine if the line is a good fit. The resulting plots from this analysis for each phase are shown in Figure 4.4. Inspecting the χ^2 values, the relationship between fit energy and true energy is consistent with a straight line. We should mention that for this plot we only include events with true energies above 8 MeV. This value was chosen because we originally only fit events with at least 40 PMT hits, which is roughly a horizontal cutoff on Figure 4.4, and a horizontal cutoff would bias events in low energy bins to be too high. By 8 MeV, there were basically no events with only 40 PMT hits. We could have made the cut higher, since we are only interested in events over 10 or 12 MeV, but it is nice to have confidence that the energy fit is good over a larger range than is necessary for our analysis.

We are also interested in whether the energy reconstruction is biased in various parts of the detector. We check this in a few ways. The first is to plot the energy bias at various positions for each phase in a 2-dimensional histogram. This is done in Figure 4.5, where Subfigures (a), (c), and (e) are in $\vec{u} \cdot \vec{r}$ and $|\vec{u} \times \vec{r}|^2$ coordinates, while (b), (d), and (f) are in z and ρ^2 coordinates. Locations outside the AV (and two bins with low statistics) are in white.

In Subfigure (f) of Figure 4.5, the positions of the NCDs are plotted, and in Subfigure (e), the equatorial extent of the NCD array is sketched in, though the NCDs extend to nearly the full radius of the AV in the z-direction. The effect of the NCDs is clearly visible in (e) and (f) as an underestimation of energy for events that occur near or behind the NCD array.

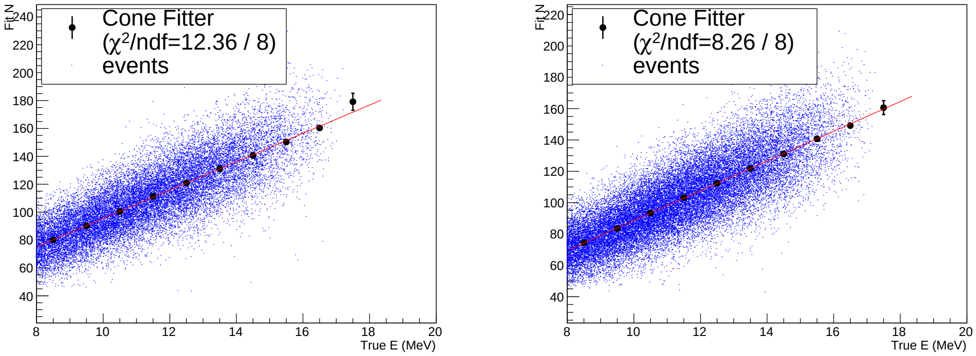

(a) Phase I *hep* energies.

(b) Phase II *hep* energies.

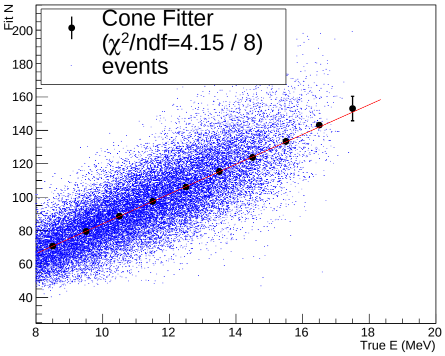

(c) Phase III *hep* energies.

Figure 4.4: Energy linearity in the three phases.

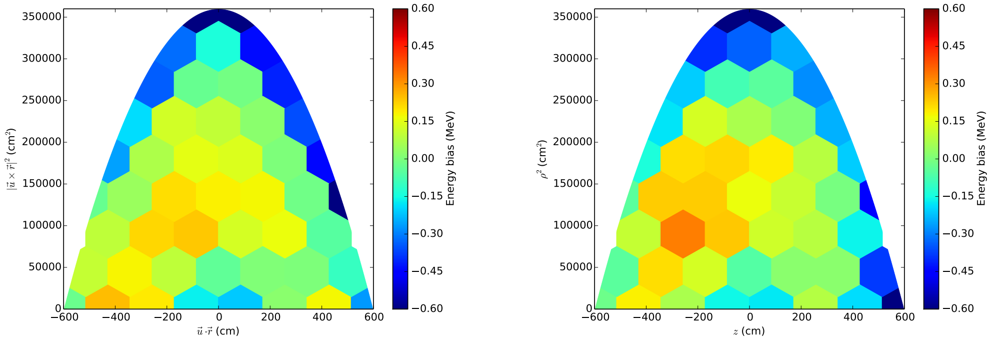

(a) Phase I energy residuals in $\vec{u}$ and $\vec{r}$ coordinates.

(b) Phase I energy residuals in z and ρ coordinates.

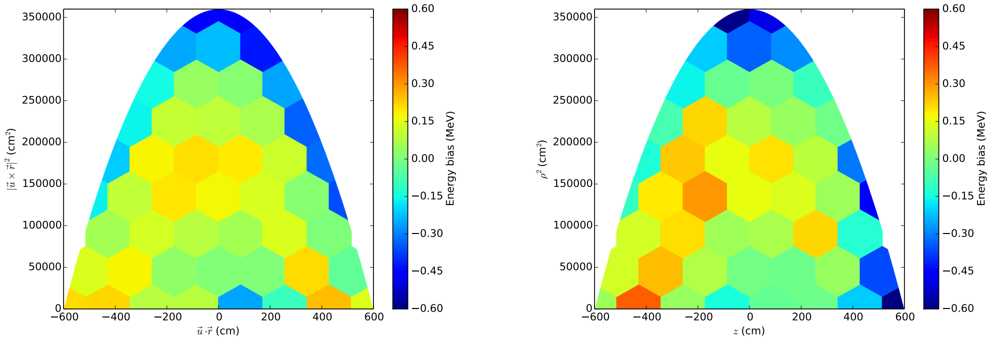

(c) Phase II energy residuals in $\vec{u}$ and $\vec{r}$ coordinates.

(d) Phase II energy residuals in z and ρ coordinates.

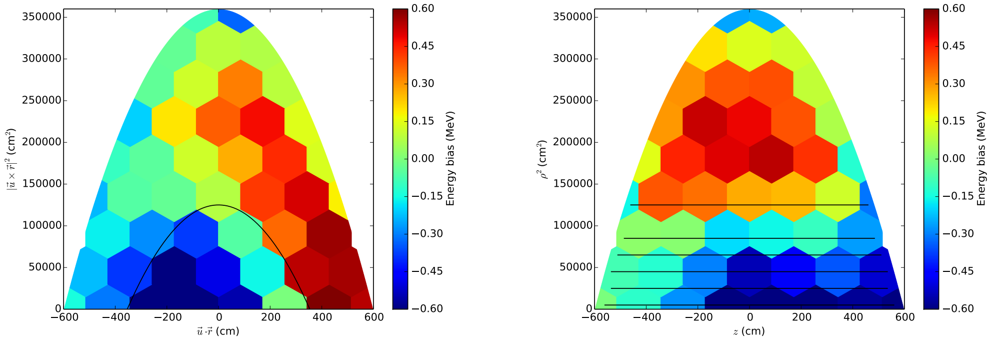

(e) Phase III mean energy residuals in $\vec{u}$ and $\vec{r}$ coordinates. Line is extent of NCDs.

(f) Phase III energy residuals in z and ρ coordinates. Lines are location of NCDs.

Figure 4.5: Cone Fitter energy bias as a function of event location.

This makes sense as we have made no correction for shadowing of the PMTs by the NCDs. Events at high $\vec{u} \cdot \vec{r}$ or high $|\vec{u} \times \vec{r}|^2$ are in front of or beside the NCDs, and since they are not obstructed by the NCDs, they are estimated to have slightly higher energy. This makes sense as the average energy over all the events is calibrated to be correct, but events that are behind the NCDs have a bias to slightly lower energies.

A few other patterns are discernible. Events at very high radius and near the center of the detector seem to fit to a slightly lower energy, while those elsewhere in the detector seem to fit slightly high. FTK has a similar effect which is attributed to the PMT angular response function [19] and corrected after-the-fact. In our case, the underestimate at high ρ specifically could be partly due to the increased absorption and scattering in the hold-up ropes and belly plates. Events very close to the neck also have energies that are underestimated. This could be because events with a lot of light in the neck are vetoed, so there is a bias toward lower energy events in this bin.

In Phases I and II, the above biases are fairly small (under 0.5 MeV), and even though they are larger in Phase III, they are still less than 1 MeV which is still not larger than the uncertainty in the energy given the number of photons detected. Additionally, we will see these biases are smaller than those in other fitters.

Another way to search for bias in Cone Fitter, and to compare it to FTK and RSP, is to bin events by their location in the detector, specifically, z and ρ^2 (cylindrical coordinates), and $\vec{u} \cdot \vec{r}$ and $|\vec{u} \times \vec{r}|^2$ (coordinates relative to the electron direction). A totally unbiased energy fitter would have a mean residual consistent with zero in each bin. It is also informative to measure the fitters' precision by plotting the standard deviation of residuals in each bin. In Figures 4.6 through 4.9, we show the mean residuals binned by coordinate (Subfigures (a), (c), and (e)) as well as the spreads in the residuals (Subfigures (b), (d), and (f)), for each phase and for each coordinate mentioned. These results show Cone Fitter, RSP, FTK (including all events), and FTK with cuts applied. In each plot, the error bars are one standard error (on the mean and standard deviation, as appropriate).

The cuts applied to FTK are those used in previous analyses[38][28], and include several parameters. The parameter β_{14} measures isotropy, so it is helpful for eliminating non-electron events. The in-time ratio, ITR, allows us to cut events with too much late light. The parameters phiprb and a2dprb are figures of merit for position fitters (in this case FTP), that measure whether an event has a Cherenkov-like distribution. Finally, FTK returns a positive and negative uncertainty, which are also used as figure-of-merit cuts. Several cuts depend on energy, and in the energy range relevant to *hep* events the cuts are tighter because there is less spread in β_{14} and ITR for these higher-energy events. Because the electron events are meant to fall within these cuts, and we are analyzing Monte Carlo electron events, the change in cut values does not significantly affect which of the higher-energy events are cut.

For events with $E_{\mathrm{FTK}} < 14$ MeV (the ^{8}B range), the cuts are:

$$\begin{aligned} -0.12 < \beta_{14} < 0.95 \\ 0.55 < \mathrm{ITR} \end{aligned} \tag{4.1}$$

and for events with energy greater than 14 MeV (the *hep* energy range), the cuts are:

$$\begin{aligned} 0.2 < \beta_{14} < 0.8 \\ 0.65 < \mathrm{ITR} \end{aligned} \tag{4.2}$$

and for both energy ranges, the following cuts are applied:

$$\begin{aligned} r_{\mathrm{FTP}} &< 550 \text{ cm} \\ \mathrm{phiprb} &> 10^{-5} \\ \mathrm{a2dprb} \times \mathrm{Nhits}^{6} &> 10^{5} \\ \text{FTK positive uncertainty} &< 1.195/\sigma_{\mathrm{FTK}} \\ \text{FTK negative uncertainty} &> -0.972/\sigma_{\mathrm{FTK}} \end{aligned} \tag{4.3}$$

where

$$\sigma_{\mathrm{FTK}} = -0.0919 + 0.353\sqrt{E_{\mathrm{FTK}}} + 0.0278E_{\mathrm{FTK}} \tag{4.4}$$

In most cases the fitter that appears to have the largest bias at extreme coordinates and the largest spread at all coordinates is FTK without cuts. Meanwhile, the fitter that

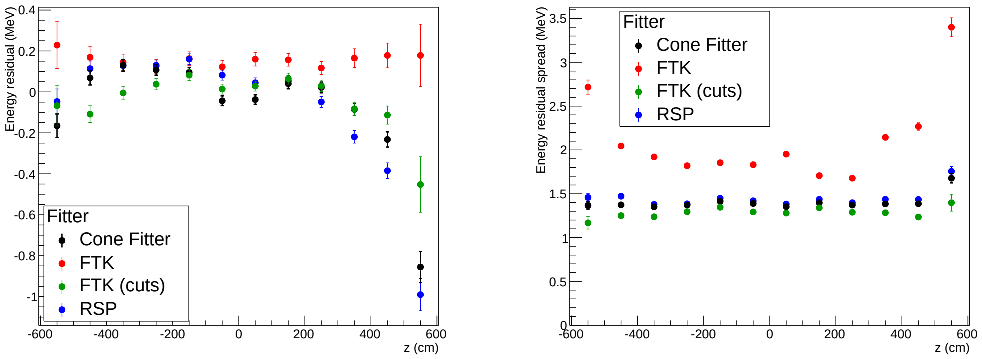

(a) Phase I mean energy residuals.

(b) Phase I energy spread.

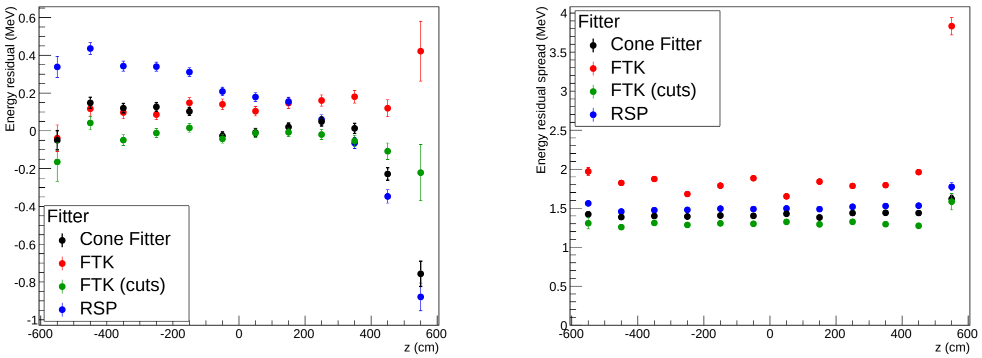

(c) Phase II mean energy residuals.

(d) Phase II energy spread.

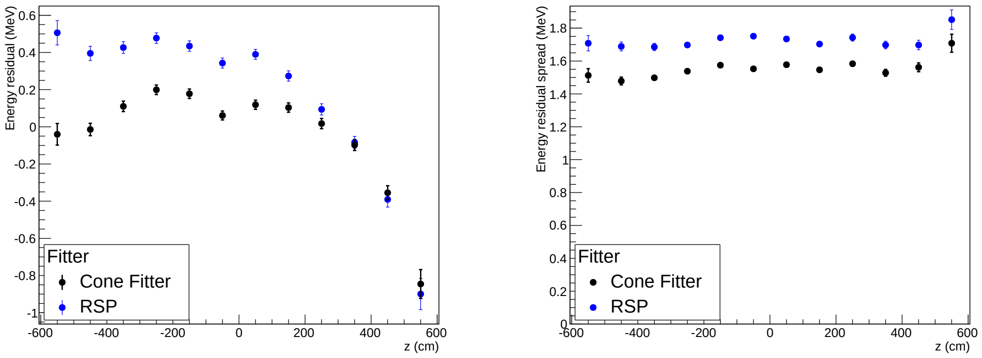

(e) Phase III mean energy residuals.

(f) Phase III energy spread.

Figure 4.6: Energy performance as a function of z-coordinate.

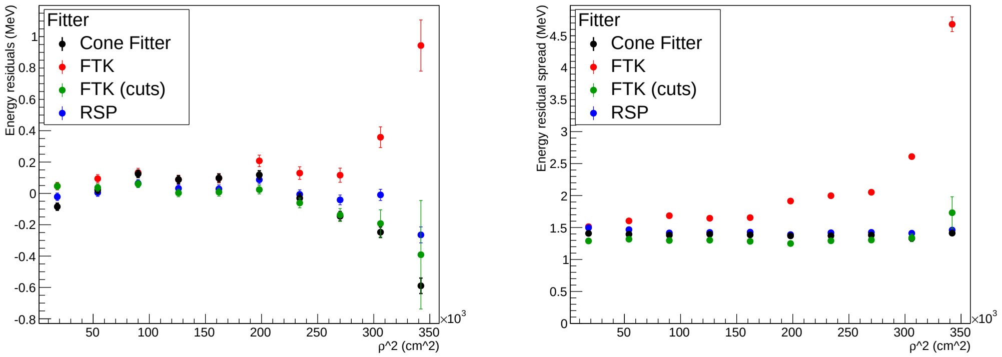

(a) Phase I mean energy residuals.

(b) Phase I energy spread.

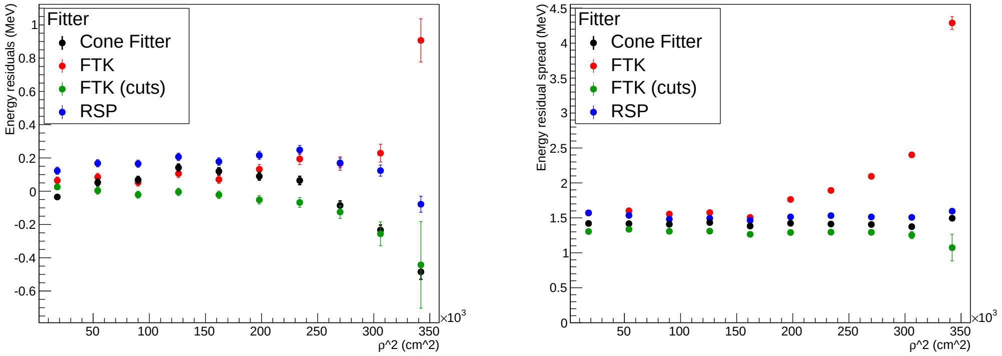

(c) Phase II mean energy residuals.

(d) Phase II energy spread.

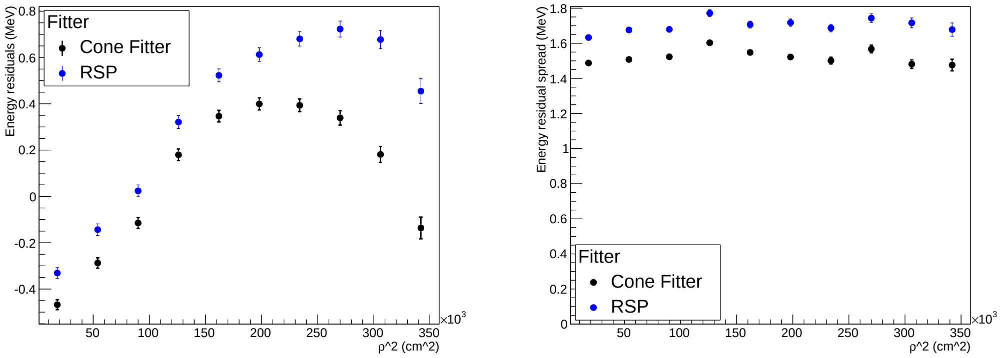

(e) Phase III mean energy residuals.

(f) Phase III energy spread.

Figure 4.7: Energy performance as a function of ρ-coordinate.

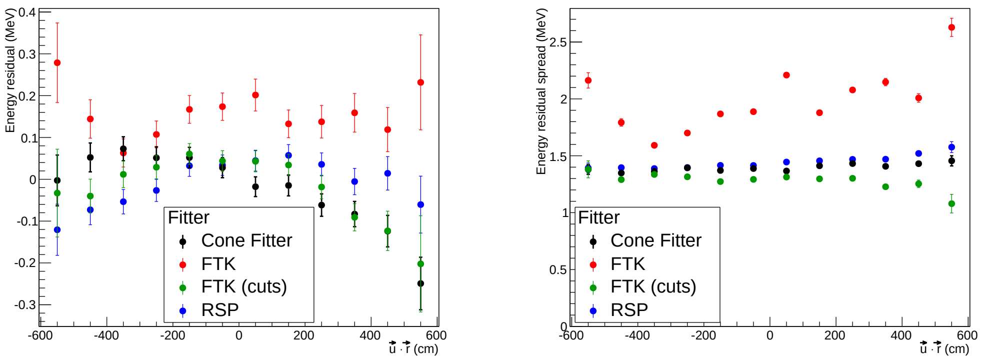

(a) Phase I mean energy residuals.

(b) Phase I energy spread.

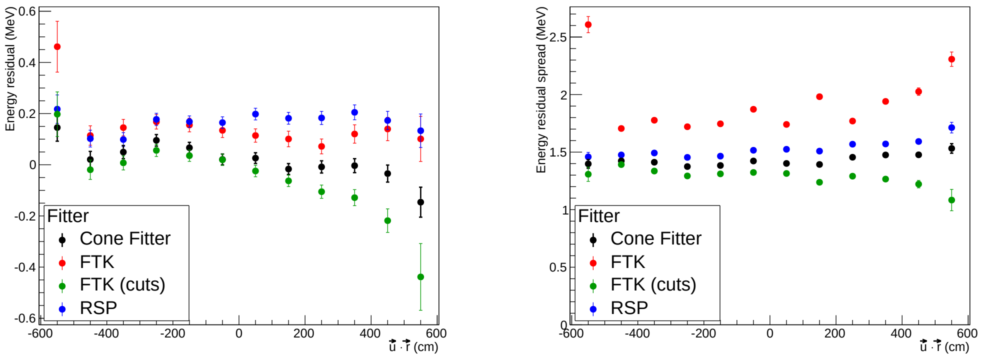

(c) Phase II mean energy residuals.

(d) Phase II energy spread.

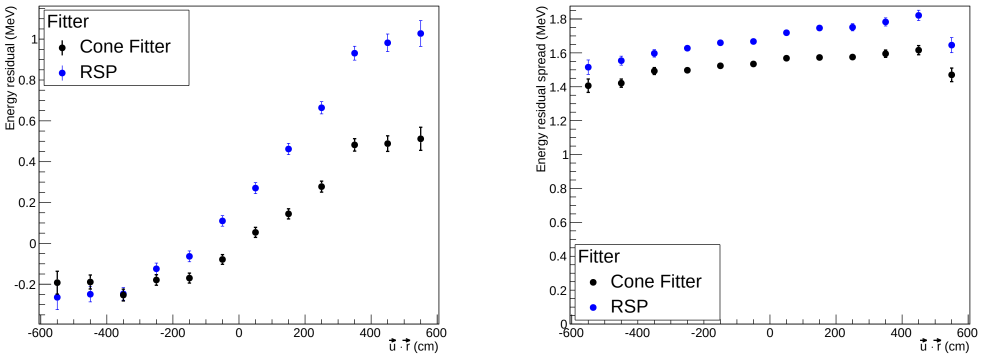

(e) Phase III mean energy residuals.

(f) Phase III energy spread.

Figure 4.8: Energy performance as a function of $\vec{u} \cdot \vec{r}$.

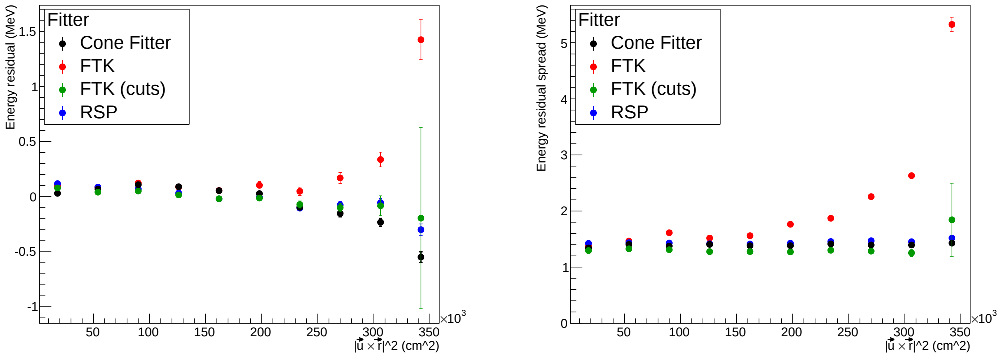

(a) Phase I mean energy residuals. (b) Phase I energy spread.

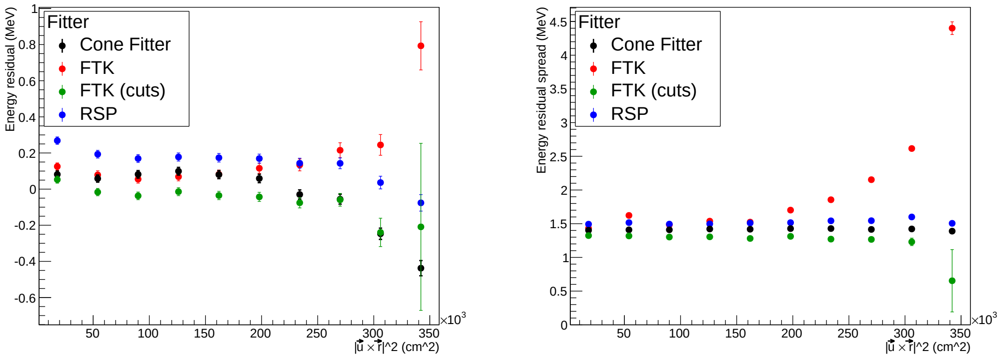

(c) Phase II mean energy residuals. (d) Phase II energy spread.

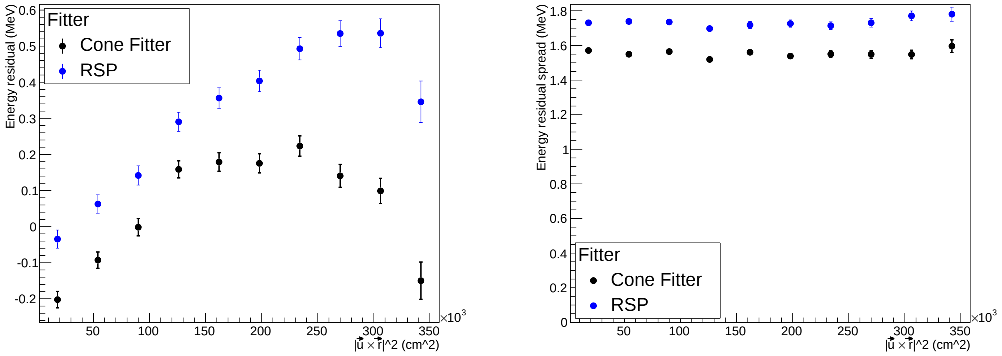

(e) Phase III mean energy residuals. (f) Phase III energy spread.

Figure 4.9: Energy performance as a function of $|\vec{u} \times \vec{r}|^2$.

usually appears to have the least bias and smallest spread is FTK with cuts. These two observations can be explained by a few comments about FTK and the cuts applied to it. First, FTK's overestimation of energy for events within the bulk of the acrylic vessel means that events at extreme positions will be unreliably fit by FTK. One of the cuts applied is a radius cut, so these events are not included in the "with cuts" plots. Second, FTK estimates the uncertainty in the energy it returns for each event, and another cut specifically removes events with high uncertainty. This will also tighten up the spread and reduce the bias. Finally, FTK uses an ad hoc correction to energy based on the radius [19], while Cone Fitter has no such correction.

Cone Fitter (without cuts) tends to perform only slightly worse than FTK with cuts, but with the advantage of not needing to cut events. Since a goal of this work is to increase the amount of data for the *hep* analysis in part by pushing the fiducial volume to the entire D_2O region, reliably reconstructing energy at these positions is a big win. Note that the volume of a 6-m-radius sphere is 30% larger than a 5.5-m-radius sphere, so this is a significant gain.

To estimate the effect of potentially applying an ad hoc radial correction to Cone Fitter, like the one applied to FTK, I tried binning the events from Phase I into twenty bins in r^3 (so that there were roughly the same number of events in each bin), and then finding a best-fit line between Cone Fitter's energy parameter and the true energy for each bin separately. Having calibrated the events in this way, I then repeated the analysis that created Figures 4.6 through 4.9, and added up the squared means from each bin. The results are shown in Table 4.1. By this measure, Cone Fitter has a smaller bias than FTK+cuts in all coordinates except z, and if z and ρ are combined, Cone Fitter has a smaller bias than FTK+cuts.

Despite the decreased bias achieved with an ad hoc radial correction, we choose not to perform one for the final analysis. The main reason is that we have such limited statistics of real events anyway, and binning them by radius would result in even fewer events in each bin. A possible solution we chose not to pursue would be to adjust the calibrated energy values in a position-dependent way using Monte Carlo data as a basis for the position-dependent correction. This option is unappealing because the low statistics make it difficult to validate

	z	ρ^2	$\vec{u} \cdot \vec{r}$	$\|\vec{u} \times \vec{r}\|^2$
Cone Fitter	0.358	0.043	0.045	0.011
FTK+cuts	0.254	0.220	0.077	0.073

Table 4.1: Combined bias of fitters in given coordinates, calculated as the sum of squared mean residuals from each bin.

such a correction.

A final observation from Figues 4.6 through 4.9 is that Cone Fitter has smaller biases and spreads than RSP in nearly every case. This is important because in Phase III RSP was otherwise the best energy fitter available.

Though we have not done so, some aspects of the fitter performance could be validated with data. For example, the ^{16}N calibration source produces γs that Compton scatter electrons. Fitting data from this source and comparing to Monte Carlo would give confirmation of the fitter energy resolution, at least in the neighborhood of 6 MeV. However, the Monte Carlo data for this source was incomplete so we were unable to fit the Monte Carlo for the comparison. There is also some complication with using this source as γs are not precisely the same as electrons, so we do expect some mismatch between them. Another potential source of data validation would be to fit Michel electrons. Michel electrons occur over a variety of energies and are readily identified by their coincidence with a muon, though they occur preferentially at energies near 50 MeV, above our region of interest. Also, the statistics are quite low (only a few hundred events per phase).

4.4 Acrylic Vessel Instrumental Backgrounds

Another goal of this work is to reliably distinguish Acrylic Vessel Instrumental Backgrounds (AVIBs) from ordinary neutrino events. The technique we developed for doing so was to fit the event with both a one-cone distribution as well as a model specially developed for AVIBs, and produce a cut based on the likelihood ratio of these two fits. Specifically, find the mean likelihood in each case, take the ratio of the two likelihoods, and then take the logarithm of this value. For short, we will call this the log-likelihood-ratio. The fit procedure is described in more detail in Section 3.7.1.

To test this technique, we use Monte Carlo data created with a variety of assumptions. The events are assumed to be monochromatic, with wavelengths of either 335, 400, or 500 nm and are created either on the surfaces or in the bulk of the acrylic. Since the physical source of these events is not known, it is also not known which of the wavelengths or source locations is most realistic. That said, previous comparisons with data indicate that 335 nm light created on the inner surface of the acrylic vessel is most similar to real AVIBs [38].

In addition to simulated AVIBs, the test described was also performed on Phase I Monte Carlo *hep* events and also on a selection of real data (a 1/3 blinded data set from Phase I, cut between 80 and 120 PMT hits, which is a little below the energy region of interest for *hep* analysis. There was no radius cut because this is fitter-dependent). The *hep* events are so we can estimate the sacrifice from the cut to remove AVIBs, and the real data are to show how real events are distributed into electron and AVIB types.

Previously, the best parameter to identify AVIBs was found to be β_{14}. Figure 4.10 (a) shows β_{14} as a function of energy for simulated AVIBs and electron events. There is pretty considerable overlap, and the distribution of β_{14} values appears to remain fairly uniform as a function of energy. Even though β_{14} does a reasonably good job of finding AVIBs, it is not reliable enough to be used to cut AVIBs. This is one of the reasons for the fiducial volume cut at $r = 550$ cm.

Figure 4.10 (b) shows the log-likelihood-ratio as a function of energy. The separation is improved, and higher-energy events have increased separation. This is likely due to the fact that detecting more photons gives the fit more information, and therefore it can be more definitive about which event type is correct. In particular, events above 15 MeV (the most important region for *hep* analysis) have nearly perfect separation.

Figure 4.10 (c), (d), and (e) sort the AVIBs by phase, wavelength, and whether they were

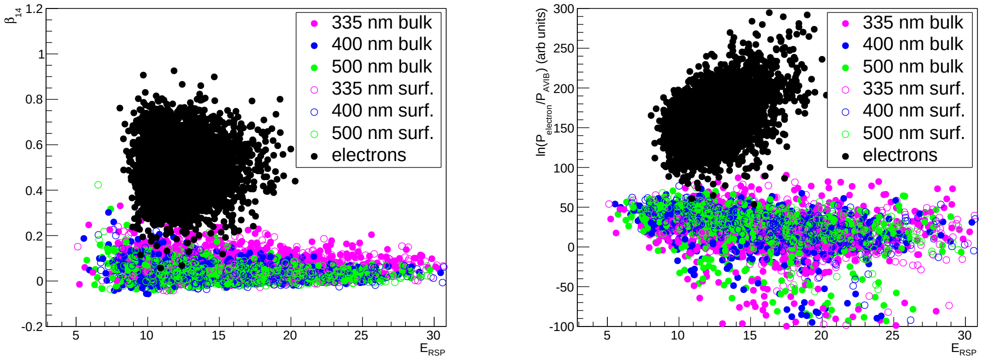

(a) β_{14}, the previous cut for AVIBs.

(b) Log-likelihood-ratio vs. energy.

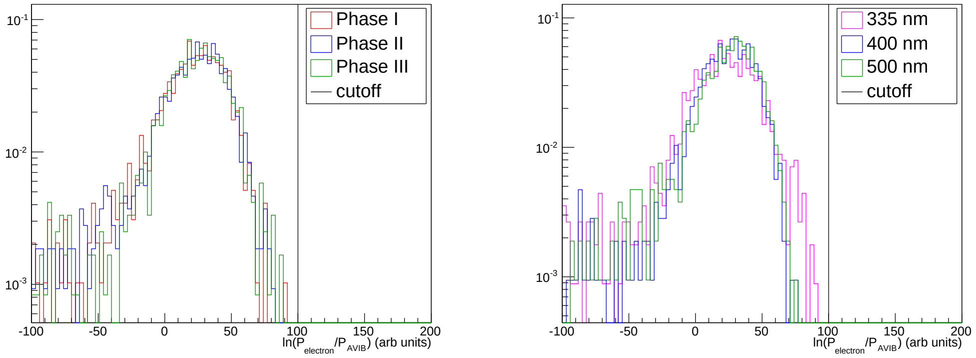

(c) AVIBs sorted by phase.

(d) AVIBs sorted by wavelength.

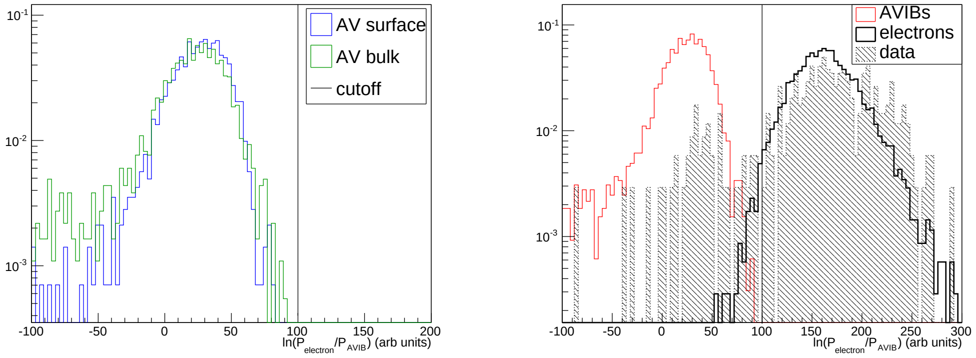

(e) AVIBs sorted by location.

(f) Comparison between MC and data

Figure 4.10: Cutting AVIBs

generated on the surface or in the bulk of the AV. These distributions all look extremely similar, which gives us some confidence that the exact details of the Monte Carlo are not critical for a cut on the log-likelihood-ratio to effectively remove the AVIBs. A cutoff value of 100 cuts every AVIB event simulated, so we expect no contamination from AVIB events.

Figure 4.10 (f) shows a comparison between the MC events (both AVIB and *hep*) and real events. There are clearly some AVIB events in the real data. There are more real events at high log-likelihood-ratio than the Monte Carlo *hep* events would indicate. This is easily explained by the fact that the real events do not have a radius cut, so these events are probably low-radius events that look very unlike AVIBs. Otherwise, the real data looks very similar to the electron data plus some AVIBs.

Despite the evidence that a wide variety of potential AVIB events are effectively removed by this cut, it is possible that real events might thwart our general models. We can estimate potential contamination from AVIBS that do not look like any of our models. To do this, we can shift the distribution of the likelihood ratio of our AVIBs to improve the match between data and Monte Carlo. The data has a peak at a log-likelihood-ratio value of 40.6, while the 335 nm, inner surface MC has is peaked at 26.0. If we shift the MC distribution by 14.6 units, a total of 3 out of 5912 MC AVIB events, or 0.05%, now pass the cut. Since there appear to be 50 AVIB events in the Phase I, 1/3 data (by counting the number of events below the cutoff), we can expect about 0.1 AVIBs to pass the cut in Phase I, or about 0.3 AVIBs over the full three phases.

To estimate the sacrifice from this cut, we can examine the the Monte Carlo *hep* events. Of 3480 events tested, 46 are below this cutoff. This is a 1.3% sacrifice, which is actually an overestimate because we are most interested in higher-energy events, and the lower-energy events are most likely to be cut. For example, with a 12 MeV cut, the sacrifice is reduced to 1.1%. Also, the MC *hep* events in the test all were created with a radius greater than 550 cm, since this region is most similar to AVIBs. Assuming events at lower radii are cut more easily, the other 77% of the events will have a smaller sacrifice, possibly much smaller. Given these considerations, 1.0% sacrifice is a reasonable estimate, and 2.0% is very conservative. Since this sacrifice is smaller than other sources of uncertainty we expect in the final result, we do not worry about attempting to decrease them further.

Chapter 5

RUN SELECTION

The SNO experiment collected data for six years, from 1999 through 2006. During that time, there were three phases characterized by different strategies to detect neutrons from neutral current events. There were approximately 1364 days of run time that was designated as "neutrino" runs, from which the 1073-day[1] "golden" analysis runlist was chosen. Another 261 days of run time were tagged by operators as "unusual conditions" (UC), some 30 days of which were included in the golden analysis run set. Another 534 days that the detector was operational were spent in calibration, either with a source or calibrating the electronics or PMTs. There were also a variety of other states that made up the small remaining fraction of run time.

There are different challenges to detecting a *hep* signal in SNO data than there are to measuring the ^{8}B signal. A critical part of the ^{8}B measurement, for example, was distinguishing elastic scattering, charged current, and neutral current events. However, *hep* events in the 15-18 MeV region will be nearly all charged current events (with very a small contribution from elastic scattering). Also, lower-energy radioactive backgrounds overlap with the ^{8}B spectrum, but ^{8}B itself is the largest background for *hep* detection. Finally, the rate of *hep* events is expected to be very low, so we are attempting to detect them or place a limit on their rate, whereas the ^{8}B event rate was sufficiently high to measure it precisely. Therefore, it is reasonable to use different criteria to select runs for this *hep* analysis than were used for the ^{8}B analysis.

We expect the runs that were accepted for ^{8}B analysis to also be acceptable for *hep*

---

[1] Run times in this work are "real" time data was being taken. Other works use "live" time, which is corrected for dead time, so the times differ slightly between this and other works.

analysis. We also expect that some additional runs that were rejected from ^{8}B analysis for low-energy considerations will also be acceptable for the higher-energy *hep* analysis. While it would be easier to simply use the ^{8}B run list, we think it is worthwhile to use the available data as completely as possible.

The largest challenge to a *hep* measurement is the low event rate, so using more of the detector's live time is a straightforward way to increase statistics, as long as we do not disproportionately increase backgrounds as well. This effort is complementary to discussion in Chapter 3 where we increase the fiducial volume by suppressing AVIBs, the only high-energy background localized at high radius.

5.1 Run lists

Our task in this chapter is to identify specific events, runs, or categories of runs that we can add to our data set. There are roughly three types of data that we examine: 1) data used in SNO ^{8}B analyses, 2) superficially normal runs that were excluded from previous analyses, and 3) calibration runs. Data of type 1, that which was used in previous analyses, is considered for the purposes of run selection to be golden data, and will be used to compare with the other types of data. Data of types 2 and 3 are our test data.

Runs of type 2 were not considered in previous analyses because of concerns like work near the detector, PMT malfunctions, anomalies in the shift report, and others. However, it is not known if there is actually anything anomalous about this data (or if possible anomalies are in the region of interest for *hep* analysis) so some fraction of it may be useful for study of *hep* neutrinos.

Type 3 data is calibration data. Several low-energy sources were inserted into the detector in a variety of configurations to calibrate various aspects of the detector. Due to an abundance of caution, a significant portion of running time was devoted to these calibrations. We may be able to add these runs to our data set without adding backgrounds in the region of interest. The higher event rate due to calibration sources may also affect the run time we use, but we will attempt to correct for this.

We generated a total of 24 lists of runs, four golden run lists and 20 test run lists:

1. D_2O golden runs used in 2011 paper [5]
2. D_2O runs from 2006 paper [3] excluded from 2011 paper for high radon
3. D_2O neutrino runs not used in either paper
4. D_2O ^{16}N calibration runs
5. Salt golden runs used in 2011 paper
6. Salt runs previously identified as suitable for *hep* analysis with high radon
7. Salt runs previously identified as suitable for *hep* analysis with activated ^{24}Na
8. Salt runs previously identified as suitable for *hep* analysis not fitting into either above category
9. Salt runs that were excluded because they were early in salt phase
10. Salt neutrino runs excluded from ^{8}B analysis and existing *hep* list
11. Salt neutrino runs with UC bit set
12. Salt ^{16}N calibration runs
13. Salt Th calibration runs
14. Salt Cf calibration runs
15. D_2O-2 runs (a period after salt phase before NCD phase) previously identified as suitable for *hep* analysis and which have the D_2O circulation off
16. D_2O-2 runs previously identified as suitable for *hep* analysis and which have the D_2O circulation on
17. NCD golden runs used in 2011 paper with D_2O circulation off
18. NCD golden runs used in 2011 paper with D_2O circulation on
19. NCD neutrino runs excluded from 2011 ^{8}B analysis
20. NCD runs with a radon spike
21. NCD ^{24}Na calibration runs
22. NCD Th calibration runs
23. NCD Cf calibration runs
24. NCD AmBe calibration runs

To create these lists, I first tabulated the amount of run time from runs longer than 30 minutes. This cutoff was used because shorter runs are missing most of the processing that was applied to longer runs. Runs that were included also needed to be the specified run type ("neutrino" runs or runs of the specified calibration source type, but not transition runs, source-moving runs, supernova runs, etc.). With the exception of golden runs and runs in list 11 (which is specifically a "UC" run list) and 8, 15, and 16 (which were previously identified as good), the UC bit must not have been set. Included runs also must have the D_2O circulation off (except for golden lists and lists 8 and 16). NCD phase runs were only considered if they occurred after the designated commissioning phase, which ended at run 48219. Run lists were also checked to ensure that data is available for all runs included.

The golden run lists (1, 5, 17, and 18) include some runs with the UC bit set, with the D_2O circulation on, and a few anomalous cases. The NCD phase golden run list had enough runs with the circulation on that I separated these into a different list for comparison. In general, I tried to only include data with a single non-ideal property (e.g. normal in every way except the "UC" bit was set, or the D_2O circulation was on, or it was a calibration run, but not two of the above). I also was only interested in types of runs that amount to a considerable increase in run time.

To be clear about expectations, I predict that lists 2, 6, 7, 8, 15, 16, and 20 are very likely to look exactly like the golden run lists (1, 5, 17, and 18) because radon produces lower energy decay products and/or because these runs have been previously identified for *hep* analysis. I predict that at least some of the runs on lists 3, 9, 10, 11, and 19 are likely *not* usable, since they were probably examined in detail already and rejected. I predict that neutron calibration runs (13, 14, 22, 23, 24) could be usable because neutrons in the first two phases produce lower-energy γs, and in the third phase capture in the NCDs producing no light. It would not be surprising if the neutron source runs in the salt phase look worse than in the NCD or D_2O phase, since the γs emitted are at a higher energy and may therefore have a tail in our energy region of interest. Sodium runs (list 21) produce high-energy gammas which may show up in our energy region of interest. Nitrogen-16 decays (lists 4 and 12) produce low-energy γs, which is promising, but they occur at a high rate so pileup can be a problem and the high rate causes them to fail low-level (DAMN) cuts.

5.2 Kolmogorov-Smirnov test

There is not really a standard way to compare multivariate data sets and determine whether they are consistent. However, the accepted way to compare 1-dimensional data sets is with a Kolmogorov-Smirnov test ("KS test"). The KS test is widely used because it works equally well with theoretical or experimental data, it works well with data sets of very different sizes, and doesn't require a parametric function for comparison. Unfortunately, it translates poorly to more than one dimension because it operates on ordered distributions and the problem of ordering multi-dimensional data is poorly-defined. There have been attempts to develop multi-dimensional KS tests, but none are as elegant as the 1-dimensional KS test, as they use procedures like imposing a large number of arbitrary orders on the data sets and then doing statistics on the set of results.

A possibility we briefly considered was performing 1-dimensional KS tests on a variety of parameters (for example, energy, isotropy, radial position of the event, z-position of the event, etc.). However, many of the obvious parameters are not independent, so a slightly unlikely distribution may appear very unlikely since it will show up in more than one parameter. Also, even for truly independent parameters, if we consider many of them, we expect to get a few unlikely results just by chance, so we risk rejecting a lot of good data this way.

5.3 Wald-Wolfowitz test

A less well-known test for comparing 1-dimensional data sets is the Wald-Wolfowitz test (WW test) [21]. The algorithm to compare two 1-dimensional data sets is as follows. First, combine the data sets, but keep track of which set each point belongs to. Second, arrange the combined data set into order. Third, split the combined data set into sequences of adjacent data points originating from the same set. Finally, count how many such sequences are in your combined set. This number is the test statistic, and we can convert it to a p-value by finding the probability that random assignment of the points to two sets would result in the same (or worse) test statistic.

Let us work through a simple example. Set A contains the values $\{61, 59, 55, 50, 54\}$ and set B contains the values $\{42, 54.5, 51, 47, 53\}$. (Perhaps these are measured heights of trees in two different parks, and you wish to know whether the trees have the same distribution of heights.) First we combine the data sets and order them: set AB contains $\{42, 47, 50, 51, 53, 54, 54.5, 55, 59, 61\}$. We now replace the members of set AB with labels corresponding to their original set, preserving the order: $\{B, B, A, B, B, A, B, A, A, A\}$. We now partition this into same-origin sequences: BB, A, BB, A, B, AAA. We have 6 such sequences, so our test statistic is 6. The probability of having exactly six sequences is given by combinatorics: of all the possible orderings of 5 As and 5 Bs, how many have 6 sequences? There are $\frac{(N_A+N_B)!}{N_A!N_B!} = 252$ (equally probable) ways to order 5 As and 5 Bs. Of these, 72 contain six sequences.

Therefore we find that $P(6) = 72/252 = 28.5\%$. To make this into a p-value, we need to know the probability of getting a result *at least as bad* as 6. In this case, doing worse means having fewer sequences (after all, for very different sets, we would get a result like $AAAAABBBBB$, giving a test statistic of 2, and the best-possible overlap looks like $ABABABABAB$, giving a test statistic of 10). Table 5.1 gives the number of permutations (exact and cumulative) and p-values associated with each test statistic.

So in our example, with a p-value of 0.64, we cannot rule out these two sequences coming from the same distribution, but a case with, for example, only 2 or 3 sequences would have a p-value below 0.05 ($p < 0.05$ is a typical, but arbitrary, cutoff for rejecting the null hypothesis).

The WW test, it turns out, underperforms the KS test in one dimension, which is why it is less popular than the KS test. However, it translates to higher dimensions much more naturally, and does an increasingly satisfactory job the higher the dimensionality of the data. The reason is that while *order* is poorly defined in multiple dimensions, *distance* is actually a quite natural description of the data. The scheme to extend the WW test to multiple

Sequences	Exact Permutations	Cumulative Permutations	p-value
2	2	2	0.008
3	8	10	0.04
4	32	42	0.16
5	48	90	0.36
6	72	162	0.64
7	48	210	0.83
8	32	242	0.96
9	8	250	0.99
10	2	252	1.0

Table 5.1: Wald-Wolfowitz 1-dimensional example p-value calculation

dimensions is to find graphs (that is, sets of connections) of points that are close together, and then count the number of cuts between subgraphs whose points originate entirely from the same input data set. In our analysis, this proceeds as follows. First, combine the two data sets to be compared (i.e. the golden data and the test data), keeping track of each point's original data set. Then, calculate the distance between each pair of points[2]. From this array of distances, construct an optimal graph connecting the points. Then, split the graph into subgraphs of adjacent points from the same original data set. Finally, count the number of cuts between subgraphs, which is the test statistic, and calculate a p-value from it.

There are many options for the optimal graph. One choice is to use each point's nearest neighbor. As suggested by [21], we employ the minimal spanning tree. A tree is a graph that connects a set of vertices with exactly one possible path between any given pair of points (which is to say there are no loops). A spanning tree connects all points in the set, leaving

---

[2] We will discuss how to calculate this distance later.

no orphans. A minimal spanning tree (MST) is a spanning tree such that the total length of all edges is the minimum possible. Several algorithms exist to calculate the MST for a data set. We use the Kruskal algorithm, which begins with each point as a separate tree, then successively searches for the smallest connection between two trees and joins the trees via this connection. Once all trees have been joined into one tree, the result is the MST.

Let's work through a 2-dimensional example of the WW test. Set C contains the five ordered pairs

$$\{(61,15),\ (59,24),\ (55,7),\ (50,26),\ (54,3)\}$$

and set D contains

$$\{(42,9),\ (54.5,17),\ (51,32),\ (47,15),\ (53,1)\}.$$

One could imagine C and D are the positions of trees in two parks, and one wishes to know whether their locations were determined by the same natural process. As before, we combine the sets into CD. We then construct the MST, which is shown in Figure 5.1a. Now, we cut the MST into subgraphs of points from the same set. In this example, we make five subgraphs by making four cuts, which is our test statistic (the four dotted lines in Figure 5.1a). To convert this to a p-value, we must again calculate the probability of requiring at most four cuts if we randomly assigned the vertices in our MST to sets C or D. We take the shape of the tree as given.

There are still 252 ways to assign ten points to two sets of five. We again find every possible assignment of the points to the two sets, and count the number of cuts in each case, and calculate a p-value from the probability of getting no more than four cuts. This calculation is shown in Table 5.2.

In this example, we calculate a p-value of 0.37 and cannot rule out these two distributions being the same. Inspection of the table reveals that it is almost not possible to rule out the distributions being the same (at $p < 0.05$), as the worst possible p-value is 0.047. The possibility of the test being too weak to rule out the null hypothesis is a concern with the real data as well, since we wish to be able to compare a small set (perhaps a single data

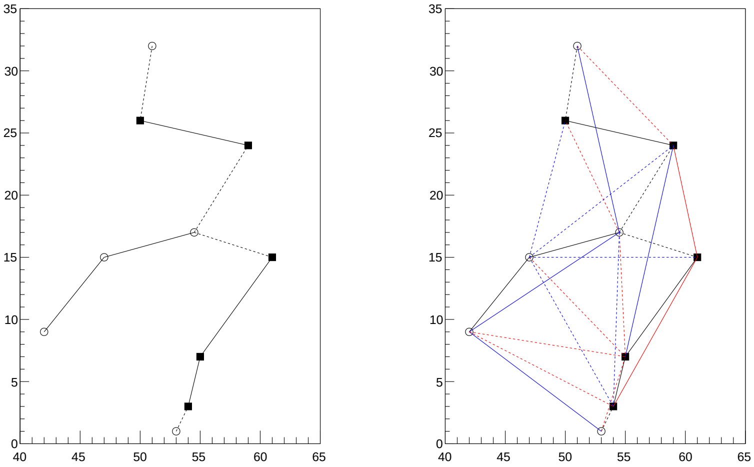

a) b)

Figure 5.1: a) Minimal Spanning Tree for the example. Solid lines are within subgraphs and dotted lines are “cut”. The number of cut lines is the test statistic. b) 1st (black), 2nd (red), and 3rd (blue) MSTs for the same example.

point) to a large set. This leads to possible numbers of cuts being quite limited (perhaps only 1, 2, or 3), which causes the possible calculated p-values to take on only a small number of values. A solution proposed in reference [21] is to calculate a graph that contains many more edges than the MST, such as the union of 1st through mth MSTs. To calculate a 2nd MST, first calculate the MST (which we now call the 1st MST), then delete all the edges used in the 1st MST from the list of possible edges. Calculate a new MST from the remaining edges, which is called the 2nd MST. By repeating this process m times, m MSTs can be calculated, which can then be analyzed as one combined graph. Our test statistic is still the number of cuts between points with different original data sets. Having sampled the data in this improved way, we typically get a more satisfactory distribution of possible WW test parameters, giving us a useful statistical test.

For testing sets with a large number of entries in each, reference [21] provides analytical formulas for the mean and standard deviation of the (approximately normally-distributed)

Cuts	Exact Permutations	Cumulative Permutations	p-value
2	12	12	0.047
3	26	38	0.15
4	56	94	0.37
5	64	158	0.63
6	56	214	0.85
7	28	242	0.96
8	8	250	0.99
9	2	252	1.00

Table 5.2: Two-dimensional Wald-Wolfowitz example p-value calculation.

test statistic. In cases where one of the sets is a single event, it is straightforward to calculate the p-value exactly by trying each possible assignment of events to the sets. In cases with only a few events in one of the sets, the large-N limit is inappropriate, but it is still computationally intensive to consider every possibility, so we settle for a random sampling of possible assignments to estimate the p-value.

In the above 2-dimensional example, the notion of "distance" was straightforward. However, with the SNO data we wish to use values like radius, energy, isotropy, and others that live in totally different spaces, and we would get nonsense if we ignored units of the various parameters. We need a scheme to put these parameters on roughly equal footing. We choose to define distance as the Mahalanobis distance, which is calculated as follows. First, treat each data point as a vector of the various parameters to be tested, even if the various parameters have very different units and magnitudes. Then, calculate a covariance matrix C for each parameter, where

$$C_{ij} = \frac{1}{N} \sum_{k} (X_{ik} - \bar{X}_i)(X_{jk} - \bar{X}_j) \tag{5.1}$$

for N data points X, where k is over all data points, i and j are over the coordinate indices, and $\bar{X}_i$ indicates the mean value of coordinate i. The covariance matrix allows us to scale each parameter by its standard deviation and simultaneously adjust the weighting of distances by the correlation between parameters. We then invert the covariance matrix, and then calculate the distance between a pair of data points $\vec{u}$ and $\vec{v}$ with the formula

$$d_M = \sqrt{(\vec{u}-\vec{v})^T C^{-1} (\vec{u}-\vec{v})} \tag{5.2}$$

The Mahalanobis distance corrects for dissimilar parameter magnitudes, handles correlation between various parameters elegantly, and will work with data of any number of parameters.

5.4 Testing the test

Since the modified WW test is not widely used, it seems appropriate to test that it works as expected before using it for run selection. For our golden data set, we used 100 Monte Carlo ^{8}B CC events. For our test data sets, we used sets of 5 additional Monte Carlo ^{8}B CC events, Monte Carlo neutron events, and Monte Carlo AVIBs, as sets we expect to pass, fail, and fail, respectively.

For realistic data, we needed to decide which variables to put into the WW test. We had many variables to choose from, so we picked a set that represented position, time distribution, isotropy, and energy. We tried to avoid variable combinations that were basically duplicates because that would increase weight on the duplicated variable and also increase the weight of discrepancies between the variables, which we do not necessarily expect to be meaningful.

For position, we decided to use x, y, and z. We expect events to be uniform in Cartesian coordinates, so effects of a calibration source should be obvious. While it is tempting to use r to identify events near the AV, r is correlated with the other variables (e.g. a high z event is necessarily at high r) and would make events in the center of the detector most likely to look like outliers, rather than high-radius events.

For the time distribution variable, we used in-time ratio (ITR). ITR is the fraction of light that is detected by a PMT within a small window of the expected arrival time of light at that PMT (in other words it rejects early and late light). This was the standard choice for previous SNO analyses.

SNO used two variables to measure the isotropy of light in an event, θ_{ij} and β_{14}. While θ_{ij} is conceptually simpler (it is just a measure of the average pairwise angular separation of PMT hits), β_{14} (a linear combination of spherical harmonics of the PMT hit pattern) proved better at discriminating background events, so we used β_{14}.

Finally, for our energy variable we chose to use the RSP energy. We chose RSP over FTK because FTK is known to reconstruct the wrong energy for events it believes are inside the AV. We chose not to use Cone Fitter energy for the practical reason that Cone Fitter is computationally intensive and we wanted to proceed with run selection before Cone Fitter fits were complete. We also considered using the number of hit PMTs, since this is the simplest measure of energy, but this varied over time as PMTs failed, so we settled on RSP.

We also had freedom to choose the number of MSTs to calculate. If too few MSTs are used, the resulting distribution of calculated p-values has few values, is not flat, and loses the power to produce p-values below 0.05. If too many trees are used, the calculation times become long and the marginal trees do not improve the test statistic (in fact, in the extreme limit of too many trees, nearly all the possible connections are used and the test becomes very weak again). According to [21], this last criterion is that the number of trees must be much less than half the number of data points. Results with ten trees are presented here, though the results are similar with other numbers of trees.

Figure 5.2 shows the results of the test. All subfigures use the same 100 ^{8}B CC reference events. Subfigure (a) shows sets of 5 ^{8}B CC test events, and as expected the p-values are uniformly distributed. Subfigure (b) shows sets of 5 neutron test events. Neutron events differ subtly from electron events, which is demonstrated by the preponderance of low-p-value tests. Subfigure (c) shows sets of 5 AVIB test events, which differ significantly from the ^{8}B electron events. The WW test appears to work as expected.

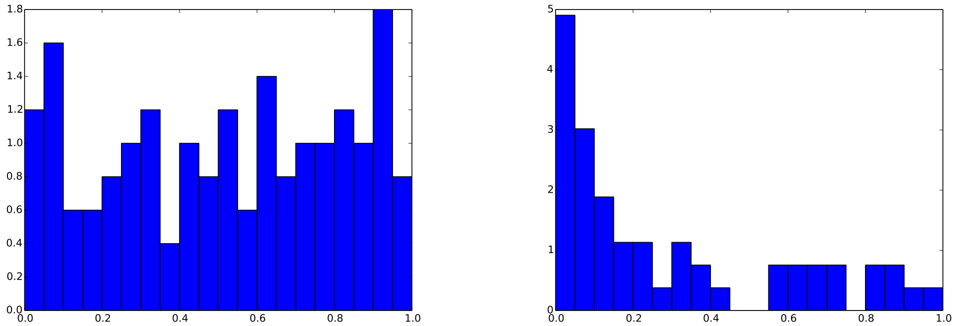

(a) p-value for test data with ^{8}B CC events.

(b) p-value for test data with neutron events.

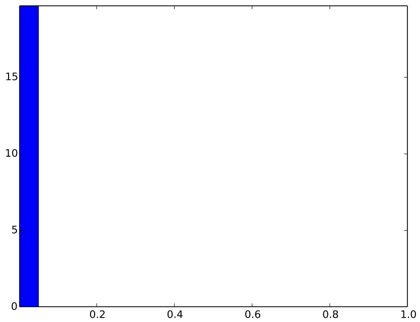

(c) p-value for test data with AVIB events.

Figure 5.2: Results of a series of tests of the WW test.

5.5 Run rejection strategy

In addition to the distributions of event parameters, we also need to check that we detect a sensible rate of events in each test set. We can do this by using Poisson statistics to rule out runs and types of runs that have an improbably high rate of events compared to the golden data sets.

Ideally, an entire run type (for example, runs using the AmBe source) would have a number of high-energy events consistent with the golden runs. Barring this, a small subset of such runs would be cut (by assumption something would be wrong with these runs other than the presence of the calibration source), and the remaining runs would be consistent with the golden runs.

Given that an entire test set or a large fraction thereof has a reasonable event rate, we can then apply the WW test to compare the test set to the golden runs. This will be done in two steps. First, each event in the test set will be compared to the golden set. Events with a p-value less than 0.05 or alternatively which return the minimum possible p-value shall be flagged. If an unreasonably high number of these are detected, the run type will be rejected.

If the individual events in a test data set appear to have a sensible distribution of p-values (that is, there are not a surfeit of low p-values), then the full test set shall be compared to the golden set with a WW test. A p-value of less than 0.05 here shall also reject the test set.

To ensure we are only comparing relevant events, we apply several cuts to the data. We apply a radius cut,

$$r < 650\ cm. \tag{5.3}$$

We chose this value because we hope to increase the fiducial volume to the full size of the AV ($r = 600$ cm), but there is some uncertainty in the position of events (at 2σ, about 30 cm). We want to include as many of these events as possible, but also cut events near the PMTs, so we settled on $r = 650$ cm.

We also apply a cut on β_{14}, the event isotropy parameter:

$$-0.12 < \beta_{14} < 0.95 \tag{5.4}$$

This is the same cut that was applied in the ^{8}B analysis to remove events whose distribution of photons does not look like a Cherenkov cone.

We also cut on the in-time ratio:

$$0.55 \leq \mathrm{ITR}, \tag{5.5}$$

which is also the cut used in the ^{8}B analysis.

We employ an energy cut:

$$12\ \mathrm{MeV} < E_{\mathrm{RSP}} < 15\ \mathrm{MeV}. \tag{5.6}$$

The energy range was chosen to be higher than expected calibration source backgrounds (especially neutrons) but below *hep* events to avoid biasing the sample of events in the *hep* energy range.

Some of the run lists to be considered do not have RSP energies (these will be noted below). None of these runs are calibration source runs, which means the energy cutoff is less critical because we do not expect increased low-energy backgrounds. In these cases, the number of hit PMTs will be substituted for energy, within the range $100 < N_{\mathrm{hits}} < 140$, which should represent approximately the same energy range.

SNO employed a series of low-level cuts called DAMN cuts [36], and we utilize the set that passes the cuts corresponding to the mask:

$$\text{DAMN mask} = \texttt{0x0CF56FE1}, \tag{5.7}$$

which is the same mask used in ^{8}B analyses.

Finally, there is a semi-random blinding applied to the data. The data is split into 30 sets, with each set being a contiguous series of events of equal time. Each set is assigned a bit, 1-30, and we select only events which are assigned one of the bits corresponding to the hexadecimal mask:

$$\text{Data set} \in \texttt{0x2aab084}. \tag{5.8}$$

Contemporary SNO analyses have all agreed to use only those events in this 1/3 of data until we agree to unblind.

5.6 Results from first test

We applied the above cuts to the set of run lists enumerated in Section 5.1. The length (in days) of the run list, the number of events that pass the cuts, and the probability of detecting at least that many events is shown in Table 5.3. The probabilities are all calculated using the total rate across all golden runs (139 events across 1073.8 days), with the exception of run lists 9, 11, 15, and 16. As these lists do not contain RSP energies, we cut on the number of PMT hits as discussed above. Additionally, because the relationship between PMT hits and energy varied over time, these latter runs were only compared to the event rate in golden salt runs (25 events across 398.2 days).

From the Poisson probabilities in Table 5.3, listed in the column headed "Rate *p*-value", we can definitely exclude run lists 14, 21, 22, 23, and 24, calibration runs which have far too many events. Although list 13 has no events, list 22 uses the same source and has too high a rate, so we will conservatively exclude list 13 as well. Run lists 4 and 12 will be excluded because they are very short. (Originally, ^{16}N runs constituted a lot of run time, but after removing runs that were less than 30 minutes, little run time remained.) List 8 has a slightly high rate, with $p < 0.05$, so it is also excluded.

Lists 7 and 21 both mention ^{24}Na, and though list 21 fails, we do not also cut list 7. The reason these are different is that list 21 contains very high event rates from ^{24}Na that was purposely added to the detector for calibration purposes. List 7, meanwhile, has a very small amount of ^{24}Na which was cosmogenically created in the salt added to the detector.

After considering only the event rate, we are left with run lists 2, 3, 6, 7, 9, 10, 11, 15, 16, 19, and 20 which may still be added to the *hep* analyses, a total of 249.6 days. We observe, however, that 8 of the 20 test lists had either 0 or 1 event, so both of the rate test and WW test are very weak and do not give convincing evidence these runs are okay. This will motivate expanding energy window in order to increase the number of events considered and make this test much more powerful.

Next we check the remaining run lists for individual runs with two or more events within

List	Days	Events	Rate p-value	List WW p-value	1-ev. WW $N(p<0.05)$	1-ev. WW $N(p_{\text{worst}})$
1. D_2O golden	283.1	30	0.88			
2. D_2O Rn	29.6	6	0.19	0.25	0	0
3. D_2O ν other	13.4	2	0.52	0.52	0	0
4. D_2O ^{16}N	2.7	0	1.00	-	-	-
5. Salt golden	398.2	53(25*)	0.44			
6. Salt *hep* Rn	50.8	8	0.34	0.08	0	1
7. Salt *hep* ^{24}Na	18.6	1	0.91	0.89	0	0
8. Salt *hep* other	18.4	6	0.03	0.60	0	0
9. Salt ν early*	21.1	2*	0.81	0.06	1	0
10. Salt ν other	26.3	1	0.81	0.09	0	1
11. Salt ν UC*	31.5	4*	0.58	0.06	0	0
12. Salt ^{16}N	5.2	0	1.00	-	-	-
13. Salt Th	16.2	0	1.00	-	-	-
14. Salt Cf	23.5	9	0.00	0.00	0	0
15. D_2O-2 no circ*	8.6	3*	0.10	0.42	0	1
16. D_2O-2 circ*	4.6	0*	1.00	-	-	-
17. NCD golden	354.3	55	0.10			
18. NCD gdn. circ	38.2	1	0.99			
19. NCD ν other	9.3	1	0.70	0.70	0	0
20. NCD Rn	16.4	0	1.00	-	-	-
21. NCD ^{24}Na	14.4	37	0.00	0.00	0	0
22. NCD Th	12.8	39	0.00	0.00	1	6
23. NCD Cf	13.3	188	0.00	0.00	0	5
24. NCD AmBe	29.1	11	0.00	0.00	0	1

* Used number of PMT hits because RSP energy was not available.

Table 5.3: Results of run selection tests

our cuts. The golden salt list, 5, contains a run (23221) that contains two events; it was 24 hours long, which makes about a 1% probability of having two events. Since 147 salt runs were over 20 hours, we expect about one to have two events. Run list 15 contains a 7-hour run with two events (34691). There is a 0.08% chance any given 7-hour run would have two events within our cuts. Since 2275 runs across all phases were 6-8 hours long, so we also expect one of these to have two events. After examining the shift reports for the two runs with two events, nothing suspicious was found so they will be included.

The fifth column in Table 5.3, headed "List WW p-value", shows the result of the WW test comparing each test list to its respective golden run list. Every calibration run we eliminated based on event rate also fails the WW test, which is reassuring. As further evidence that the WW test works, originally a few salt UC runs were inadvertently included that had the "bubbler" turned on (the details of which are not important). Three events were located in roughly the same position at the bottom of the AV, and the WW test returned a p-value of 0.0007. After removing those runs, the p-value became 100 times better, 0.06.

Examining the table, list 8 fails the rate cut but looks fine according to the WW test. Four lists just barely pass the WW test: lists 6, 9, 10 and 11. List 6, the *hep* runs with high radon, and list 9, the early salt-phase neutrino runs, have p-values of 0.08 and 0.06. List 10 consists of the neutrino runs that did not make it onto the *hep* or ^{8}B lists, and the WW test returned $p = 0.09$ for the single event in run list 10, which happens to be the lowest value the test could return in this case. List 11, the salt-phase neutrino runs with the UC bit set, just barely pass the WW test with $p = 0.06$. While these lists all pass our pre-determined $p < 0.05$ criterion, it is a little disturbing to have a preponderance of low p-values. Increased statistics would give us greater confidence that these runs are fine, especially lists 9 and 10 which have only 2 and 1 events, respectively.

The second-to-last column in Table 5.3 shows the counts of individual events that fail the WW test when compared to their respective golden run lists. There are only two events that have WW p-values below 0.05: one in the early salt run list, and one in the NCD thorium calibration list. The final column shows how many events returned p-values that were over

0.05 but that nevertheless were the worst possible p-value in that test. There are several of these, but it is unclear if we should consider these to be bad events or simply to conclude the test failed in these cases.

The event-by-event WW tests appear to have limited utility in identifying bad events. We can make sense of this by thinking about how the WW test works. If the test data set has only one event, every connection to that event in the graph will be a “cut”. Each point in an individual MST will usually only connect to one, two, or three other points (but conceivably more), and is much more likely to have two connections than any other possibility. The only way to get a low p-value is to have improbably few cuts, which means that the test event would need to connect to only one other event in several of the MSTs, which probably means the event is an outlier. However, as we put cuts on every variable, there were no outliers. This is clear in, for example, list 23, where there were 188 events, but none of them had $p < 0.05$ even though we expect by chance about 9 or 10 events to have had such a p-value. In the next section we address this shortcoming by increasing the number of events considered.

5.7 Expansion of the energy region

When planning the run selection scheme, we chose the energy window of 12-15 MeV because we didn’t want to bias ourselves with potential *hep* events in the 15-18 MeV range, and in order to avoid lower-energy backgrounds, particularly neutrons. While we expect neutrons to produce γs at 8 MeV in the salt phase (and lower energy in the other phases), there is enough uncertainty on the energy fit that some events could fit to 10 or 11 MeV, but we estimated that a negligible number of neutrons would appear to be as high as 12 MeV. As it happens, our first run selection test (described in the previous section) ruled out all the neutron calibration source runs anyway: ^{24}Na, Cf, Th, and AmBe. In runs that did not contain a neutron source, we expect neutron events to derive from neutral current interactions or natural backgrounds, which is to say that these should occur at the same rate in golden runs as in our test runs. Therefore, if we lower the threshold for including events in the

WW test, we can still eliminate most of the neutrons, and neutrons that pass the threshold should only cause the WW test to fail if they enter at a different rate between golden and test runs.

The energy window was chosen conservatively to avoid neutrons, and the neutron sources have been excluded. Therefore, since we need to increase our sensitivity, we decided to broaden the energy window and run the test again. This is necessary since many lists had too few events to confidently rule on whether they were similar to the golden lists or not. In particular, having run this test, we are left with five lists that we can neither rule out nor confidently include because they had either zero or one event each. Even after we expand to the full, unblinded run list, we may still reasonably expect to have an ambiguous result in these lists.

We also noticed that some run lists seem to have unexpectedly few events. Originally, we intended to eliminate runs with too many events (presumably due to a background), but we realized we may also end up with too few events if one of our cuts (for example, the DAMN cuts, which among other things remove events in close time coincidence) removes events before we see an event count. Specifically, we saw that in the ^{16}N source runs, the event rate was so high that every event was cut by the event rate, so even if there were real neutrino events, they would all be cut. Therefore very low rates are also suspicious and we will also cut events with a very low rate rate. In keeping with the standard of $p < 0.05$, we will symmetrize this and cut events with a rate corresponding to $p < 0.025$ or $p > 0.975$.

The new test is this. First, we will apply the same cuts as before, but with an energy window of 10-15 MeV, which (from the golden run lists), I expect to increase the number of events by about a factor of five. (Note that this increase in statistics is large enough as to make the new test practically independent of the old test). For the remaining lists with no RSP energy, we will again use the number of hit PMTs, and expand the window to be 90-140 hits. Then we will count events in each remaining run list, and if the rate is (by Poisson statistics) in the bottom 2.5% or top 2.5%, that list will be cut. We will then apply the WW test to the remaining lists, comparing the whole list to the golden list for its phase.

Lists for which the WW test returns $p < 0.05$ will be cut.

5.7.1 Results of expanded energy window test

The results of the expanded energy window test are shown in Table 5.4. All remaining lists pass the WW test, except list 16, the D_2O-2 list with circulation on, which we couldn't test because it still had zero events (incidentally, it is the shortest list). List 8, the salt *hep* runs that had something unusual other than sodium activation or high radon, fails the rate test with $p = 0.007$. Note this list also failed the rate test with the smaller energy window, but as it was close we tried it with more statistics. List 10, the salt-phase neutrino runs excluded from the golden list, fails the rate test with $p = 0.996$.

In all, we cut two more lists and approve nine lists to combine with the golden runs for the *hep* analysis. The additional runs amount to a combined 199.4 days of run time, an 18% increase to the 1073.8 days of golden runs. The rate of events in the golden lists is 725 in 1073.8 days, or 0.67 events/day. The test lists we approve for adding to the data set account for 199.4 days, so we expect 134 ± 23 events at 2σ. Our test lists which contain RSP energy information have 93 events, and the lists without RSP energy information have 17. In order to combine these, we scale the latter number by 261/82 to get 54 events, for a total of 147 in 199.4 days. This is consistent with the expected rate. Since the expected rate could allow up to 23 events to be included above the average rate without causing the test to fail, we treat this as an estimate of the worst-case level of background contamination from adding in the test runs. Having 23 background events out of a total of 872 is a contamination rate of 2.3% in the ^{8}B range, so this worst-case scenario is smaller than the expected uncertainty on the rate of ^{8}B events.

We can also run a WW test on the complete set of test runs that we have accepted, compared to the golden run lists. We must leave out the three lists with no RSP energy, but we can run a WW test on these vs. the salt golden runs separately. For the with-RSP test, we get a p-value of 0.86. For the no-RSP test, we get $p = 0.08$, which also passes.

List	Time (d)	Events	$p(n \geq N)$	WW-test	Pass?
1. D_2O golden	283.1	193	0.32	-	Yes
2. D_2O Rn	29.6	18	0.66	0.40	Yes
3. D_2O ν other	13.4	5	0.93	0.24	Yes
5. Salt golden	398.2	261 (82*)	0.53	-	Yes
6. Salt *hep* Rn	50.8	39	0.18	0.24	Yes
7. Salt *hep* ^{24}Na	18.6	12	0.56	0.89	Yes
8. Salt *hep* other	18.4	22	0.007	0.60	No
9. Salt ν early	21.1	8*	0.13	0.46	Yes
10. Salt ν other	26.3	8	0.996	0.18	No
11. Salt ν UC	31.5	5*	0.77	0.18	Yes
15. D_2O-2 no circ	8.6	4*	0.10	0.09	Yes
16. D_2O-2 circ	4.6	0*	1.00	-	No
17. NCD golden	354.3	235	0.46	-	Yes
18. NCD golden circ	38.2	18	0.94	-	Yes
19. NCD ν other	9.3	9	0.16	0.74	Yes
20. NCD Rn	16.4	10	0.63	0.66	Yes

* Compared number of PMT hits for test because RSP energy was not available.

Table 5.4: Results of run selection tests

List	Time (d)	Events	$p(n \geq N)$	WW-test
1. D_2O golden	283.1	9	0.19	-
2. D_2O Rn	29.6	1	0.48	0.70
3. D_2O ν other	13.4	0	1.00	-
5. Salt golden	398.2	5 (4^*)	0.94	-
6. Salt *hep* Rn	50.8	0	1.00	-
7. Salt *hep* ^{24}Na	18.6	1	0.34	0.17
9. Salt ν early	21.1	0^*	1.00	-
11. Salt ν UC	31.5	0^*	1.00	-
15. D_2O-2 no circ	8.6	0^*	1.00	-
17. NCD golden	354.3	10	0.27	-
18. NCD golden circ	38.2	0	1.00	-
19. NCD ν other	9.3	0	1.00	-
20. NCD Rn	16.4	0	1.00	-

* Compared number of PMT hits for test because RSP energy was not available.

Table 5.5: Results of run selection tests in *hep* energy region

5.8 Events in the hep energy range

We now perform the rate test and WW test on each surviving list in the energy range from 15-22 MeV (still in the 1/3 data set). For the lists without energy information, we use 140-200 PMT hits. The results of this procedure are shown in Table 5.5.

Our golden run lists average 24 events in this energy range over 1073.8 days, so we expect 4.4 ± 4.2 events in the test lists. As we only see 2 events in the test lists, we see no evidence of contamination in the *hep* energy region.

Chapter 6

EXTRACTION OF THE *hep* SIGNAL

This chapter explains the final step in the *hep* analysis: extracting a *hep* signal from the posterior energy distributions of Cherenkov events in our data set. In particular, our final result will be a posterior distribution for the rate of *hep* events which can be used to produce any desired credible region.

6.1 Cherenkov backgrounds

We expect the SNO data to include Cherenkov events with several different origins. Our signal events, *hep* neutrinos, will be one type, but ^{8}B neutrinos will compose most of the events at lower energy. These neutrino sources cannot be distinguished even in principle except statistically by the distribution of event energies.

Another type of important background event is atmospheric muons and neutrinos. When a cosmic ray strikes the atmosphere, it produces a shower of charged particles. Some large fraction of these charged particles are pions, which decay into a muon and a neutrino. The muons typically decay before reaching the detector, which produces two more neutrinos. These neutrinos extend to very high energy, and if one interacts in the detector, it can produce a wide variety of particles, including unstable spallation nuclei, neutrons, muons, and electrons. When these events occur in bursts, or when high-energy muons are present, they are easily detected. However, it is possible that all we see is a single Cherenkov event of modest energy. It is quite difficult to simulate the (very small) rate with which atmospheric neutrinos will produce such events, or what their distribution might look like, but previous studies [35] have shown that the atmospheric interactions can be well-approximated by a flat distribution in our energy range.

A specific type of atmospheric event that deserves special mention is decays of $^{12}C^*$. This isotope can rarely occur as a decay product of $^{15}O^*$ and $^{15}N^*$, which are formed when a muon undergoes a collision with ^{16}O in the water and causes it to lose a neutron or proton, respectively. Decays of $^{12}C^*$ can produce a 15.1 MeV γ that will then Compton scatter (or pair produce) and create a Cherenkov event. The rate of $^{12}C^*$ is so low it has never been measured, but unfortunately is not known. The 15.1 MeV γ is of special concern because a surplus of 15.1 MeV events can mimic a *hep* signal.

However, even a 15.1 MeV γ is unlikely to produce an event that fits to 15.1 MeV. Let us consider Compton scattering first. According to reference [22], about two-thirds of 15.1 MeV γ-rays in water will Compton scatter rather than pair produce. The electron and scattered γ will share this energy, and the γ will in all likelihood Compton scatter several more times. While the energy of all the resulting electrons must add up to (approximately) 15.1 MeV, since it is spread over several electrons, each of which have a Cherenkov threshold of roughly 0.7 MeV, the apparent energy of the event will be at least two or three MeV lower.

If instead the 15.1 MeV γ pair produces, then 1 MeV is immediately lost to the electron and positron's rest mass. When the positron annihilates, the resulting photons will be of too low of energy to produce visible electrons. Therefore, accounting for the positron's Cherenkov threshold, we expect no more than about 13.4 MeV from this sort of event.

Taking the 15.1 MeV γ possibilities together, one or two events drawn from a distribution with its endpoint in the neighborhood of 13 MeV will be indistinguishable from ^{8}B events, so we will ignore $^{12}C^*$ for this *hep* analysis. However, it would be useful if a Monte Carlo were run to determine the true energy distribution of $^{12}C^*$ events.

This leaves us with only ^{8}B, *hep*, and atmospheric events with a uniform distribution in our energy region of interest. We use the SNOMAN Monte Carlo to produce spectra for ^{8}B and *hep* events, shown in Figure 6.1 (a). The MC incorporates the neutrino spectra from reference [40], the charged current cross-section (elastic scattering is negligible at the high energies we consider), kinematics, interactions on various isotopes of oxygen in the D_2O as well as contaminants, radiative corrections, day/night effects, and many other details.

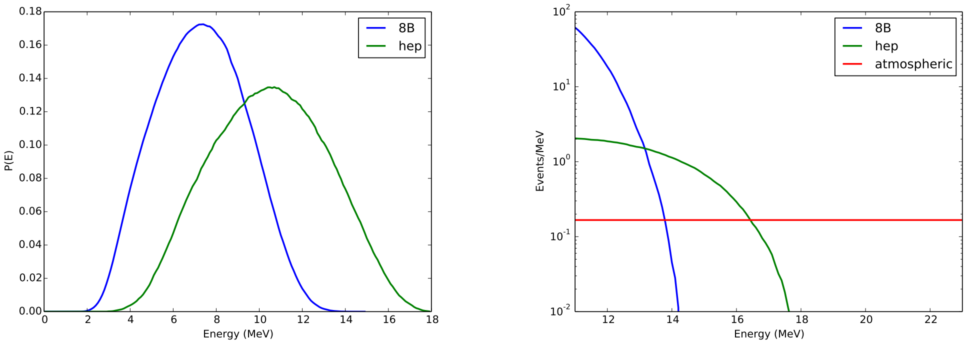

(a) Normalized spectra for *hep* and ^{8}B charged current electrons.

(b) Expected rate of events for the energy region of interest.

Figure 6.1: Expected spectra for *hep* analysis.

It does not include neutrino oscillations, but these will not affect the shape of the spectra appreciably, only the overall rate, so we can correct for oscillations when we calculate the flux.

6.2 Motivation

In order to extend the Bayesian approach of the event fitter of Chapter 3, we developed a signal extraction technique which uses the full posterior distribution of the event fitter. This is interesting partly as an academic exercise - typically one only uses the maximum likelihood value of a fitter for the signal extraction, and it is interesting to see if we can proceed without doing so - and also because with our small number of events we may expect some improvement to the fit by using as much information as we can from each event.

A specific challenge with this analysis that would not be a problem with a more traditional maximum likelihood analysis was how to apply an energy cutoff. We want to apply an energy cutoff because we are almost entirely interested in events above about 11 MeV, but even if this were not the case, we would still need to apply a cutoff because SNO itself has an energy threshold of about 5 MeV. In a standard analysis, since each event has a specific best value for energy, one simply excludes events outside any chosen energy range. However, in our case there is a *probability* any given event has an energy above or below the cutoff. Therefore, we consider both possibilities, and with hundreds of events this is computationally challenging. On the other hand, the exact location of the cutoff will have only a marginal effect on the probability any given event ends up included in the analysis.

6.3 Derivation of posterior

Ultimately we wish to report a measurement of (or limit on) the flux Φ_h of *hep* neutrinos. We can calculate the flux from the rate of *hep* events R_h in the SNO detector, the number of target deuterons in the detector N_d, and the cross-section of *hep* neutrinos with deuterons $\sigma_{h\mathrm{d}}$.

$$R_h = \int \Phi_h(E)\sigma_{h\mathrm{d}}(E)N_\mathrm{d}\,\mathrm{d}E \tag{6.1}$$

We calculate R_h assuming an event of any energy is equally likely to be detected. Later we will apply a cutoff to only consider events in a certain energy window.

As a Bayesian analysis, naturally we return a posterior:

$$P(R_h|\Xi) \tag{6.2}$$

where, as before, Ξ is shorthand for our data. In order to calculate $P(R_h|\Xi)$,[1] we will integrate over background rates for ^{8}B, R_B, and atmospheric events, R_a,

$$P(R_h|\Xi) = \iint P(R_h, R_\mathrm{B}, R_\mathrm{a}|\Xi)\,\mathrm{d}R_\mathrm{B}\mathrm{d}R_\mathrm{a}. \tag{6.3}$$

A change of variables gives

$$P(R_h, R_\mathrm{B}, R_\mathrm{a}|\Xi) = \frac{1}{R^2}P(R, \vec{f}|\Xi), \tag{6.4}$$

---

[1] To clarify the notation used in this chapter, this could also have been written $\mathrm{d}P/\mathrm{d}R_h$. As we will be working with complicated probability expressions, I find it is clearer to treat probability densities as discrete probabilities rather than write them as differentials. As we end up binning everything into histograms anyway, this is also in some ways more true to the actual calculations. To convert to differential notation, simply rewrite the probability as a differential of all variables appearing before the $|$.

where R is the total rate, $\vec{f}$ is shorthand for the fraction of the rate corresponding to each source (e.g. $f_h = R_h/R$), and the exponent of 2 on R is is the number of sources (R_h, R_{B}, and R_{a}) minus 1. Any one of the fractions is redundant, since they must add to 1.

We actually measure a number of events N rather than a rate R. Therefore, we manipulate this expression to get

$$\begin{aligned}
P(R,\vec{f}|\Xi) &= P(R|\vec{f},\Xi)P(\vec{f}|\Xi) \\
&= \sum_N P(R|N,\vec{f},\Xi)P(N|\vec{f},\Xi)P(\vec{f}|\Xi) \\
&= \sum_N P(R|N,\vec{f},\Xi)P(N,\vec{f}|\Xi) \\
&= \sum_N P(R|N)P(N,\vec{f}|\Xi) \\
&= \sum_N \frac{P(N|R)P(R)}{P(N)}P(N,\vec{f}|\Xi) \\
&= \sum_N \frac{e^{-RT}(RT)^N}{N!}\frac{P(R)}{P(N)}P(N,\vec{f}|\Xi)
\end{aligned} \tag{6.5}$$

Where the fourth line makes the substitution $P(R|N,\vec{f},\Xi) = P(R|N)$, which is justified by the observation that our only information about the true rate of events is the number of events detected, and does not depend on how many events of each type are observed. This is a direct result of the fact that the sum of two Poisson processes will follow a Poisson distribution. The result depends on $P(R)$, which is our prior knowledge of the rate, and $P(N)$, which is just a normalization factor.

We now need to find $P(N,\vec{f}|\Xi)$, which is quite similar to $P(R,\vec{f}|\Xi)$, except it has been separated into terms with known numbers of events N. N is not known ahead of time because it depends on how many events are excluded by any given cutoff. However, each event has some probability of having a true energy on either side of the cutoff, so there are many possible values of N for a data set. For now, let us take it for granted that we know the probability each event will be included in the analysis region. The details of calculating this probability will be discussed in the next section.

Let us signify an exact set of events that are included as $\{I\}$ (e.g. event 1 is included, event 2 is excluded, etc.). In that case, we can sum over all possible sets of included events:

$$\begin{aligned}
P(N,\vec{f}|\Xi) &= \sum_{\{I\}} P(\vec{f},N|\{I\},\Xi)P(\{I\}|\Xi) \\
&= \sum_{\{I\}} P(\vec{f}|N,\{I\},\Xi)P(N|\{I\},\Xi)P(\{I\}|\Xi) \\
&= \sum_{|\{I\}|=N} P(\vec{f}|\{I\},\Xi)P(\{I\}|\Xi) \\
&= \sum_{|\{I\}|=N} P(\vec{f}|\{I\},\Xi) \prod_{i\in\{I\}} P(I_i|\Xi_i) \prod_{i\notin\{I\}} P(!I_i|\Xi_i)
\end{aligned} \tag{6.6}$$

where the third line makes use of the fact that the set of included events exactly determines N for that term. The final line expands the probability of $\{I\}$ to be the product of the probability that each included event is included and each excluded event is excluded.

Finally, we need to calculate $P(\vec{f}|\{I\},\Xi)$.

$$\begin{aligned}
P(\vec{f}|\{I\},\Xi) &= \frac{P(\vec{f}|\{I\})}{P(\Xi|\{I\})} P(\Xi|\vec{f},\{I\}) \\
&= AP(\vec{f})P(\Xi|\vec{f},\{I\}) \\
&= AP(\vec{f}) \prod_{i\in\{I\}} P(\Xi_i|\vec{f}) \\
&= AP(\vec{f}) \prod_{i\in\{I\}} \int P(\Xi_i|E_i,\vec{f})P(E_i|\vec{f})\,\mathrm{d}E_i \\
&= AP(\vec{f}) \prod_{i\in\{I\}} \int P(\Xi_i|E_i,\vec{f}) \left[\sum_j P(E_i|j)P(j|f_j)\right] \mathrm{d}E_i \\
&= AP(\vec{f}) \prod_{i\in\{I\}} \int P(\Xi_i|E_i) \left[\sum_j P(E_i|j)f_j\right] \mathrm{d}E_i
\end{aligned} \tag{6.7}$$

The first line is an application of Bayes's Theorem. The second line replaces the denominator with a normalization A, and replaces $P(\vec{f}|\{I\})$ with the prior $P(\vec{f})$ because the fractional rates do not depend on which events we include (recall that $\{I\}$ is only a list of included events and does not contain information about their energies). The third line uses the fact each event is independent to rewrite the last factor as a product over included events. The fourth line rewrites the last factor as an integral over all possible values of energy E_i. The fifth line rewrites $P(E_i|\vec{f})$ as a sum over the individual probability density functions j. The final line uses the fact that fractional rates do not directly affect the data for a given event, only that energy.

Equation 6.7 leaves us with only things we know how to calculate. A is an overall normalization. $P(\vec{f})$ is the prior on the fractional rates of each type of event. $P(\Xi_i|E_i)$ can be calculated from the output of the event fitter (see the next section for a discussion). Finally, $P(E_i|j)$ is the probability density function for event type j.

6.4 Details of calculating the eventwise energy posterior

Above, we alluded to the fact that there is some probability each event has a true energy above or below a given cutoff. Now we will be careful about calculating this probability.

The probability an event is inside an energy window from E_l to E_h is given by

$$P(I_i|\Xi_i) = \int_{E_l}^{E_h} P(E_i|\Xi_i)\,\mathrm{d}E_i. \tag{6.8}$$

This expression contains $P(E_i|\Xi_i)$, while a careful reader will notice that Equation 6.7 contains the similar $P(\Xi_i|E_i)$. These two distributions are related by Bayes's Theorem:

$$P(E_i|\Xi_i) = \frac{P(\Xi_i|E_i)P(E_i)}{P(\Xi_i)} \tag{6.9}$$

We can see then that the difference between them is a the prior on energy $P(E_i)$, and an overall normalization. In Chapter 3, we assumed that the prior distribution of energy (or more specifically, the number of photons, see Section 6.6) was flat for reasons discussed in that chapter. Meanwhile, the denominator $P(\Xi_i)$ is an unknown constant that we never needed to calculate explicitly, as the MCMC algorithm automatically normalizes the posterior distribution. In the MCMC, we did not even need to be explicit about the range of energies over which $P(E_i)$ is flat, because this would simply introduce another constant factor which will be canceled by the normalization. Therefore, the expression $P(\Xi_i|E_i)$ which appears in Equation 6.7, could also be written as $P(E_i|\Xi_i)_{\text{flat prior}}$, which is exactly the output of the event fitter.

In reality, $P(E_i)$ is not flat. The most important event types, ^{8}B and *hep*, are much more likely to have energies that are lower rather than higher in the energy range of interest. Atmospheric events do not have this feature, as we assume they follow a flat distribution, but they are a very small portion of the events at the lower energy cutoff. In a model without the presence of any such backgrounds, no events could have energies above 18 MeV and vanishingly few events would be near that value.

Therefore, for any uses where we want the true posterior distribution with a realistic prior, we should calculate $P(E_i|\Xi_i)$ by using Equation 6.9, along with

$$P(E_i) = \sum_j P(E_i|j)P(j). \tag{6.10}$$

and

$$P(\Xi_i) = \int P(\Xi_i|E_i)P(E_i)\,\mathrm{d}E_i \tag{6.11}$$

where $P(j)$ is the prior any given event will be of type j. Fortunately, we expect nearly all of the events straddling the lower cutoff to be ^{8}B events, so whatever small contribution we choose to include from other distributions will have a minute effect on the probability events are included or excluded. For the upper cutoff, we expect very few events near the cutoff and we expect those events to be from the flat atmospheric distribution.

In order to avoid having to subscript posterior distributions with the priors used, we will continue to use $P(\Xi_i|E_i)$ when a flat prior is intended, and $P(E_i|\Xi_i)$ when a realistic prior is intended.

6.5 Sampling strategy

Some effort was expended attempting to calculate the contribution of every possible set of included events. This is obviously intractable in the pedestrian approach, since there are far too many terms to calculate individually. Calculating the contribution of each *event* individually, whether included or excluded, appeared promising, but this approach requires calculating the correct normalization ahead of time, which does not appear to be possible.

Therefore, the approach we settled on was to choose an ensemble of samples of events. The procedure goes as follows. First, calculate the probability each event is within the energy window. Second, select a random number between 0 and 1 for each event. Third, if the random number is less than the probability the event is in the window, assign that event to the sample. Finally, use that sample of events to calculate Equation 6.7.

As we perform this operation for a representative number of samples, we can average the resulting posteriors together to get the final posterior.

6.6 Energy scaling

A detail we glossed over above is that the fitter described in Chapter 3 returns the posterior distribution of a parameter we called N (which we will now call F for "fitter" to avoid confusion with other uses of N in this chapter) that is related to energy, but which needs to be calibrated to produce an actual energy. In Chapter 4, we showed that using data from SNOMAN Monte Carlo events, F has a linear relationship with the true energy (E) from 8 MeV to 18 MeV:

$$\begin{aligned} F &= 10.16E - 6.21 && \text{Phase I} \\ F &= 9.48E - 6.30 && \text{Phase II} \\ F &= 8.91E - 5.08 && \text{Phase III.} \end{aligned} \tag{6.12}$$

The primary strategy we will use to calibrate energy will be to float the scale and offset of the conversion in the signal extraction fit. We can incorporate prior expectations in this fit and at the end integrate over these "nuisance" parameters to produce the final posterior distribution of the rate of *hep* events.

There are two distinct sources of statistical uncertainty on the energy of a given event. First, an event that produces N_γ photons will result in a range of possible numbers of photons detected. Conversely, the number of photons detected is an imperfect measure of the total

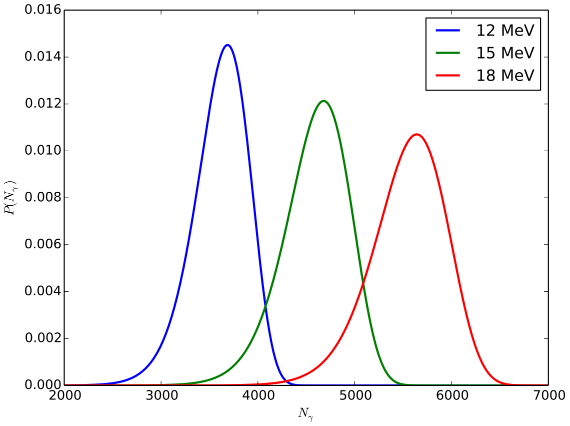

Figure 6.2: $P(N_\gamma|E)$ for three values of E.

number of photons generated by a physical event. This sort of uncertainty is already included by the posterior distribution of the event fitter, so we can imagine that $P(\Xi|F)$ is simply a scaled, shifted version of $P(\Xi|N_\gamma)$.

The second sort of statistical uncertainty is on the number of photons an event of energy E will produce. This is due to the physics of Cherenkov radiation, and as a result has a non-trivial distribution, including asymmetrical tails. The distribution $P(N_\gamma|E)$ is illustrated for a few values of E in Figure 6.2. We use the parametrization of $P(N_\gamma|E)$ from Reference [19], and observe the mean value of N_γ is linear with energy:

$$\bar{N}_\gamma = 315.3E - 142.3. \tag{6.13}$$

We must incorporate the uncertainty on N_γ into the posterior distributions explicitly.

We can do a simple test to demonstrate that the uncertainty on F and the uncertainty on N_γ, combined, account for the residuals we observe between the best value returned by the fitter and the true energy for Monte Carlo data. For this test, we use equation 6.12 to turn the mean value of F into E_{fit}, and also to convert the standard deviation to σ_{fit}, in units of energy. We also use the explicit $P(N_\gamma|E)$ to find the standard deviation σ_{N_γ}, and use this to calculate $\sigma_{\text{true}} = E_{\text{true}} \times \sigma_{N_\gamma}/N_\gamma$. We can then find

$$\frac{E_{\text{fit}} - E_{\text{true}}}{\sqrt{\sigma_{\text{fit}}^2 - \sigma_{\text{true}}^2}} \tag{6.14}$$

for each event. The mean of the resulting values is -0.027 and the standard deviation is 0.985, which are very close to the expected 0.0 and 1.0, despite the fact that we know these uncertainties are not Gaussian. This gives us confidence that we are not excluding any major source of uncertainty on the energy resolution.

As previously stated, we wish to use the detailed shapes of the posteriors to whatever extent possible. A full implementation of this might proceed as follows:

$$\begin{aligned}
P(\Xi|E) &= \int P(\Xi|E,F)P(F|E)\mathrm{d}F \\
&= \iint P(\Xi|F)P(F|N_\gamma,E)P(N_\gamma|E)\mathrm{d}F\mathrm{d}N_\gamma \\
&= \iint P(\Xi|F)P(F|N_\gamma)P(N_\gamma|E)\mathrm{d}F\mathrm{d}N_\gamma. \qquad (6.15) \\
&= \iint P(\Xi|F)\delta(AF+B-N_\gamma)P(N_\gamma|E)\mathrm{d}F\mathrm{d}N_\gamma \\
&= \int P(\Xi|F)P(N_\gamma = AF+B|E)\mathrm{d}F
\end{aligned}$$

Here we use the facts that $P(\Xi|E,F) = P(\Xi|F)$ and $P(F|N_\gamma,E) = P(F|N_\gamma)$, which is to say that F, essentially the number of photons detected, determines the data while E is superfluous, and similarly that N_γ, the number of photons produced, gives us all the information about the the number of photons detected, while again E is superfluous. We also used the fact that going from N_γ to F introduces no additional uncertainty by expressing this as a δ-function. Again, this precise relationship is not a statement that a certain number of photons produced results in a specific number of photons detected, but a reflection of the fact that the posterior distribution returned by the event fitter already contains this uncertainty.

We know $P(\Xi|F) \propto P(F|\Xi)_{\text{flat prior}}$, which is the output of the fitter, and $P(N_\gamma|E)$, which is the number of photons produced by an event with a certain energy. It turns out that this method is too computationally-intensive to use in the MCMC, because it requires us to integrate over F for every bin in E for every event, at every step.

We therefore need to simplify the approach. We observe that the parameters of $P(N_\gamma|E)$ vary slowly over the energy range of interest. For example, the standard deviation of the distribution at 12 MeV is about 8% of its mean, while at 18 MeV, it is about 7%. Therefore, incorrectly estimating the energy by an improbable margin (even a few MeV) will only result in a marginally different distribution of N_γ. We use this fact to justify estimating $P(\Xi|E)$ once at the beginning of the MCMC, creating a distribution very close to that of true energy, $P(\Xi|E')$, and then we apply small adjustments to these distributions at each step of the MCMC to produce $P(\Xi|E)$.

First, we combine Equations 6.12 and 6.13 to produce:

$$\begin{aligned} N_\gamma &= 31.056F + 50.5 && \text{Phase I} \\ N_\gamma &= 33.223F + 67.4 && \text{Phase II} \\ N_\gamma &= 35.400F + 37.5 && \text{Phase III.} \end{aligned} \tag{6.16}$$

We use these relationships to scale $P(\Xi|F)$ into $P(\Xi|N_\gamma)$. Then, we produce a histogram for $P(\Xi|E)$, and bin by bin, calculate

$$P(\Xi|E') = \int P(\Xi|N_\gamma) P(N_\gamma|E') \, \mathrm{d}N_\gamma. \tag{6.17}$$

$P(\Xi|E')$ is our best estimate of $P(\Xi|E)$. We trust the spread in this histogram to be very close to correct, but we would like to allow the exact calibration of the energy to float in the final analysis. Therefore, we parametrize the calibration

$$E = a(E' - 11 \text{ MeV}) + b + 11 \text{ MeV}, \tag{6.18}$$

where a is $\mathrm{d}E'/\mathrm{d}E$, and b is the amount by which E differs from E' at $E' = 11$ MeV. This parameterization is useful because we choose our lower cutoff to be at $E' = 11$ MeV. The corresponding value of E does not move far (as it is constrained by the ^{8}B rate). If, however, we chose to use the slope and y-intercept, they would be highly correlated in order to keep the cutoff near 11 MeV.

6.7 Algorithm

As with the event fitter, we decided to perform the signal extraction as an MCMC. For a detailed explanation of MCMC algorithms, see Chapter 3.

We begin by producing a posterior histogram of the energy parameter for each event, removing the first 10,000 steps as burn-in and using every 100th step afterwards. From these posteriors, we calculate the probability each event is between 11 and 23 MeV, using our best estimate of the conversion from fitter intensity to energy. From these probabilities, we draw an ensemble of samples of events that pass the cut.

Then, we choose a starting set of values for the parameters R_h, R_B, R_a, and for each phase, we choose a starting energy scale and offset a and b. We calculate the likelihood of our sample given the parameters, which becomes the first step in the MCMC. The likelihood (times priors) is calculated using the equation derived above:

$$\begin{aligned} P(\vec{R}|\Xi) = {} & e^{-RT}((RT)^N P(R) \\ & \times \prod_i \int P(\Xi_i|E_i) \left[\sum_j P(E_i|j) f_j\right] \mathrm{d}E_i \prod_{i\in\{I\}} P(I_i|\Xi) \prod_{i\notin\{I\}} P(!I_i|\Xi). \end{aligned} \tag{6.19}$$

The only priors we apply are on R_B and T. R_B is constrained to the expected rate, with an uncertainty corresponding to the reported uncertainty on that measurement. The prior on T contains all the uncertainties affecting the event rates in a correlated way: the run time, the fiducial volume, and the probability a solar neutrino will be detected in the electron flavor. See the next chapter for more on these uncertainties.

We then choose new values for the parameters for the next step. We calculate a new likelihood, and accept or reject the step as we did in the MCMC described in Chapter 3. After a large number of steps, the distribution of parameters should reflect the posterior distribution for that parameter. We then repeat the MCMC for a suitable number of samples (see Section 6.8) and add together their posterior distributions to find the overall posterior for the full data set.

Note that in this algorithm, the cutoffs are applied based on the fitter energy values and not the "true" energy values. We must do it this way because we must use the same sample of events for the entire MCMC (if the number of events varied, the likelihood would be impossible to interpret). Since the fitter energy values do not change step-to-step, our sample still makes sense as the energy calibration varies, we just calculate a likelihood based on varying energy cutoffs. However, as the cutoffs move in true energy units, this could cause a number of undesirable effects. For example, the cutoff will be slightly different in each phase. If the cutoffs drift to values where the probability density function for ^{8}B is dramatically different, this could also mean that when we initially chose our data sample based on the probability each event was included, these probabilities may have been incorrect. However, we do not expect these sorts of issues to be a problem in practice because the slope of the probability density functions vary slowly, and we do not expect the calibration to vary dramatically from the predicted scale and offset.

6.8 Setting number of steps and samples of the algorithm

To test the algorithm, we produced fake data from the real distributions. First, we drew approximately 30,000 (plus or minus Poisson fluctuations) energies from the ^{8}B spectrum and 30 (also with fluctuations) from the *hep* spectrum. We randomly assigned each event to a phase in proportion to the length of run time we have for each phase. We then drew a random N_γ based on the energy, and then scaled N_γ down to a value of F using Equation 6.16.

With this true value of F, we produced a histogram $P(\Xi|F)$ by first choosing a central value randomly from a Gaussian of width 10 centered on F. The value of 10 was chosen because this is approximately the average width of posteriors from Monte Carlo fits. It corresponds to roughly 1 MeV in real energy. Centered around this new value, we populated the histogram with random values following a Gaussian distribution of width 10. In this way, the probability of the true value corresponds to the weight of its bin in the histogram.

We then fit the fake data sets. We found that 10,000 steps is sufficient for burn-in. We then plotted the median, 5%, and 95% levels of the *hep* rate as a function of the step in the

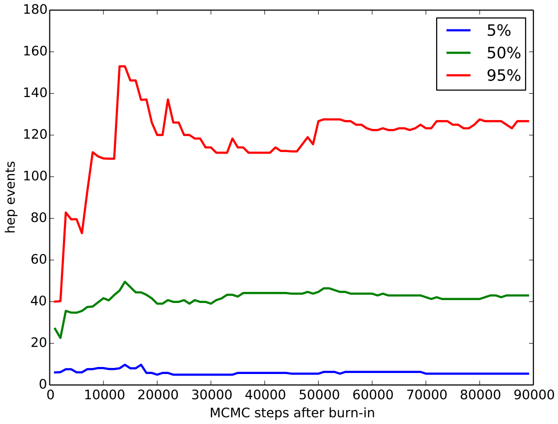

Figure 6.3: Median, 5%, and 95% levels for *hep* signal as a function of MCMC step.

MCMC. This is shown for one fit in Figure 6.3. It appears that these levels usually settle down by about 20,000 steps, but in this case there is some variation in the 95% level up to about 50,000 steps. Additionally, a calculation of autocorrelation indicates that MCMC steps are correlated for about 1000 steps or so. Therefore, 100,000 steps will be sufficient to be confident we have sampled the entire probability space.

We also plotted the median, 5%, and 95% levels for the *hep* rate as a function of number of samples drawn from the data set. This is shown in figure 6.4. From this figure, we conclude that using more than 20 samples does not change the probability of various rates appreciably.

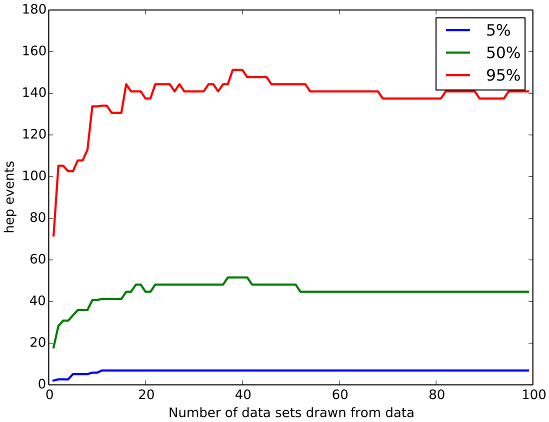

Figure 6.4: Median, 5%, and 95% levels for *hep* signal as a function number of samples drawn from the data set.

Chapter 7

ANALYSIS RESULTS

7.1 Rate prediction

First, let us calculate the expected rate of events in the SNO detector. We use the full 6-m radius volume of D_2O for this analysis. This amounts to 9.05×10^8 cm^3. The density of D_2O is 1.1056 g/cm^3, which gives 1.00×10^9 g of D_2O. At 20.0 g/mol, this is 3.02×10^{31} molecules of D_2O or 6.03×10^{31} deuterons.

The cross-section of the charged current reaction with a deuteron averaged over the ^{8}B spectrum is

$$\sigma_{\mathrm{B}} = \frac{1}{\Phi_B} \int \Phi_B(E) \sigma_{B\mathrm{d}}(E)\, \mathrm{d}E \tag{7.1}$$

which comes out to 1.177×10^{-42} cm^{-2}, while for the *hep* spectrum it is 2.958×10^{-42} cm^{-2}. The measured flux of ^{8}B neutrinos is $5.25 \pm 0.16(\mathrm{stat})^{+0.11}_{-0.13}(\mathrm{syst}) \times 10^6$ cm^{-2} s^{-1} [5]. For *hep* neutrinos, the theoretical flux is 7.97×10^3 cm^{-2} s^{-1} [11].

In Chapter 5, we calculated the livetime for our dataset to be 1273.2 days. This is only an estimate, as individual run lengths were rounded down to the next lower minute, and more importantly there are a number of sources of dead time. These sources include events occurring in close coincidence (“bursts”), coincidence with muons, and many other considerations, such as accounting for the intersections of these adjustments, which was the subject of considerable careful tabulation in previous analyses. In our increased data set, we cannot recreate this process exactly, so we remove dead time due to retriggers and bursts, and then estimate the remaining difference by assuming other sources of dead time, such as muons, occur at the same rate as in the golden data. See Table 7.1 for a breakdown of the livetime correction. We find that for the golden data set, the final livetime was 1.91% less than the time remaining after removing retriggers and bursts. When we reduce the added

Runs	Clock time (d)	Retriggers +Bursts (d)	Remaining (d)	Final (d)	Ratio
Golden Phase I	283.3	0.3	283.0	277.3	0.9799
Golden Phase II	397.4	0.2	397.2	391.4	0.9854
Golden Phase III	394.8	0.6	394.2	385.2	0.9772
Additional runs	199.4	0.6	198.8	195.0	0.9809

Table 7.1: Calculation of live time for *hep* analysis.

data sets by this fraction, we remove an additional 3.8 days, leaving us with a grand total of 1248.9 days, or 1.08×10^8 s. If we conservatively treat this 3.8 days as our uncertainty, this is still only a 0.3% uncertainty in the total live time.

Finally, due to neutrino oscillations, the survival rate of solar neutrinos in our energy window is $32 \pm 2\%$ [13]. We therefore expect 12,900 ^{8}B charged current events and 49 charged current *hep* events in the full data set, or 4300 and 16 in the 1/3 data set.

7.2 Fit with Monte Carlo data

We ran a pair of tests to simulate the one-third and full data sets using Monte Carlo rates that are close to the theoretical flux: 4300 ^{8}B events and 17 *hep* events to simulate the 1/3 data set, and 13,400 ^{8}B events and 51 *hep* events to simulate the full data set (this Monte Carlo data was created in the same way as in Section 6.8). In each case, we created and fit 100 data sets, and histogrammed the mean values of R_h, R_B, and, for each phase, the energy scales and offsets. We also show the standard deviations for all but R_h, which was very non-Gaussian. The results of this for the 1/3 data set are shown in Figure 7.1, and for the full data sets in Figure 7.2.

Let us begin by examining the rate of *hep* events. In the 1/3 data set, Figure 7.1 (a), the true value is 17 events but the means of the posteriors are around 27 events. This is easily explained by looking at the full posteriors of the rate of *hep* events, Figure 7.3 (a). Because

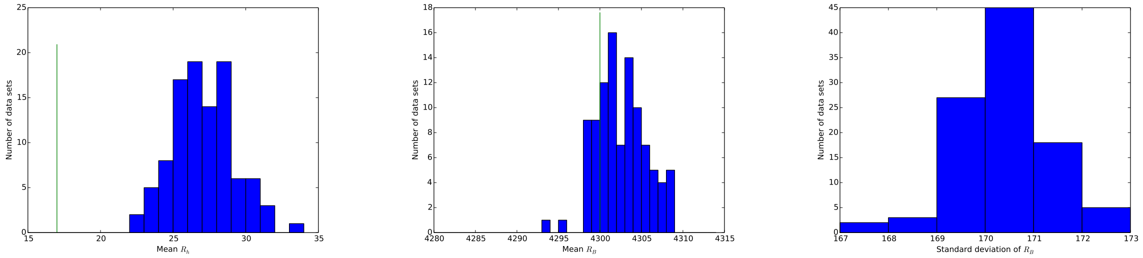

(a) Mean of *hep* rate posterior.

(b) Mean of ^{8}B rate posterior.

(c) Standard deviation of ^{8}B rate.

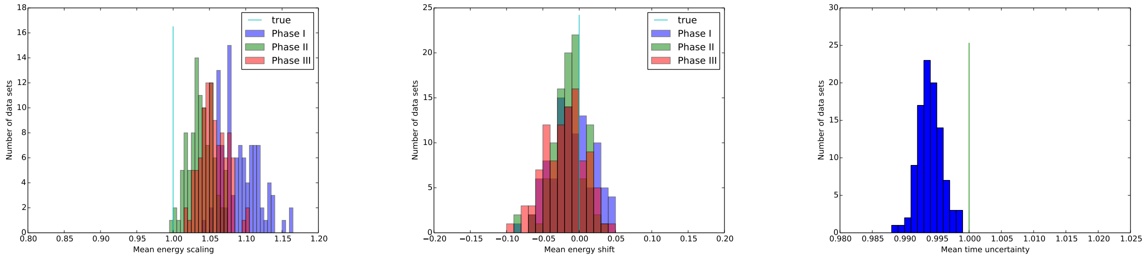

(d) Mean values of E scale.

(e) Mean values of E offset.

(f) Mean values of rate correction.

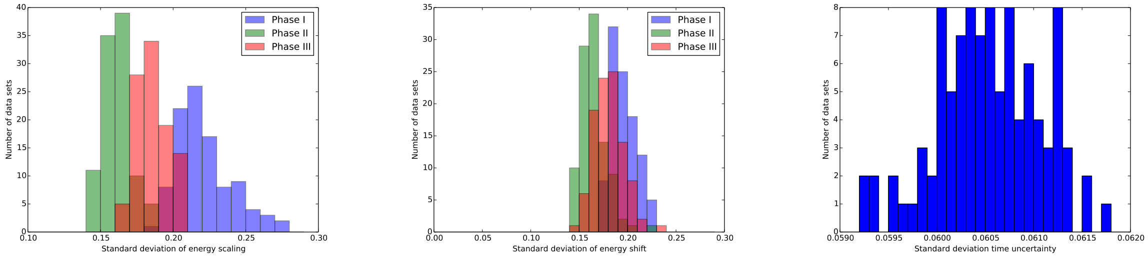

(g) St. dev. values of E scale.

(h) St. dev. values of E offset.

(i) St. dev. values of rate correction.

Figure 7.1: Summary values of the posteriors for each of 100 fake one-third data sets.

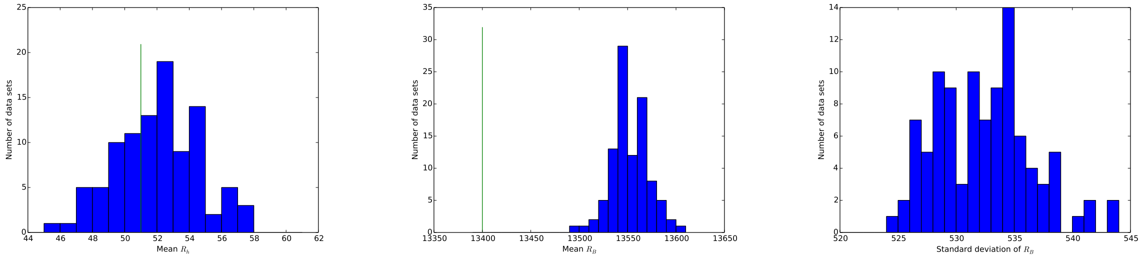

(a) Mean of *hep* rate posterior.

(b) Mean of ^{8}B rate posterior.

(c) Standard deviation of ^{8}B rate.

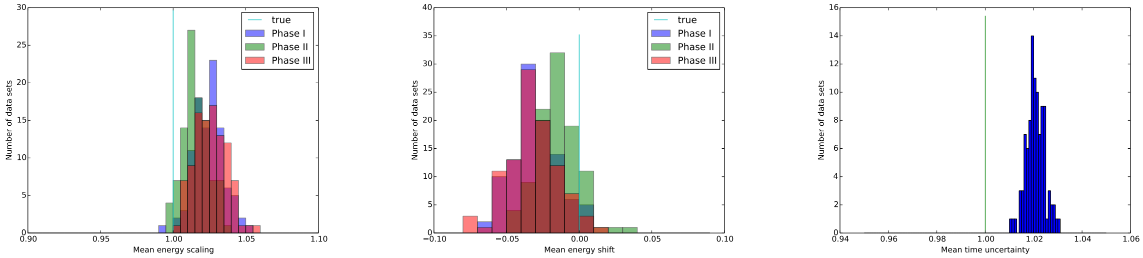

(d) Mean values of E scale.

(e) Mean values of E offset.

(f) Mean values of rate correction.

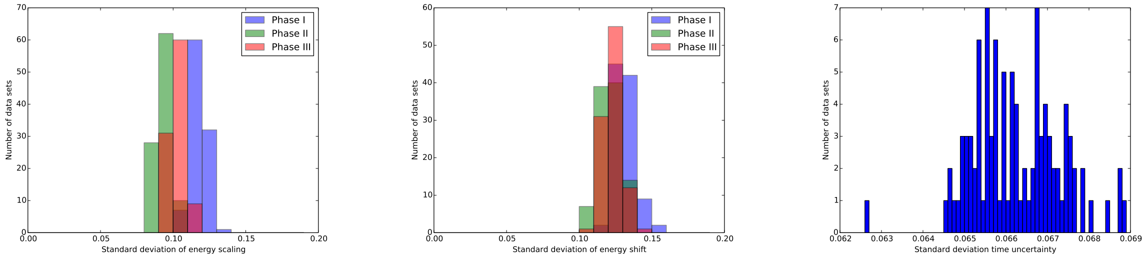

(g) St. dev. values of E scale.

(h) St. dev. values of E offset.

(i) St. dev. values of rate correction.

Figure 7.2: Summary values of the posteriors for each of 100 fake full data sets.

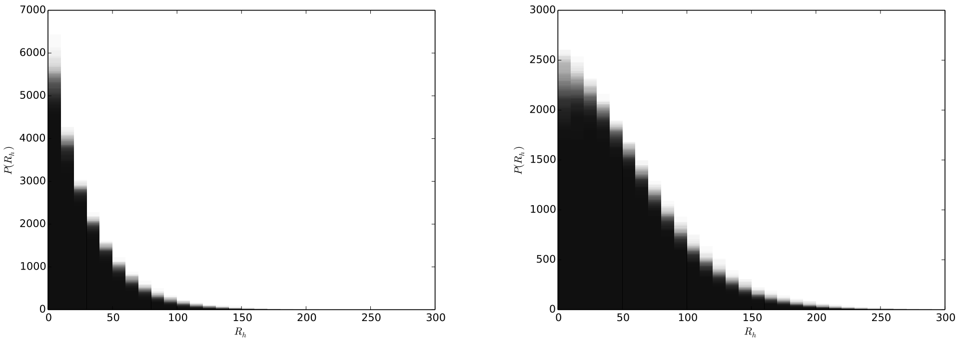

(a) One-third fake data sets.

(b) Full fake data sets.

Figure 7.3: Full posteriors in the rates of *hep* events for 100 data sets, superimposed on the same plot.

the true value is so close to zero and the rate can't be negative, the mean value is higher than the true value. In the full data sets, Figure 7.2 (a), the true rate of *hep* events is 51, and the mean values of the posteriors is around 50.

Now let us consider the rate of ^{8}B events, Figure 7.1 (b) and 7.2 (b). We constrain the ^{8}B rate in the fits with a prior of 3.8% (the uncertainty in the measured rate). Although this constraint is derived from the same data set we use, the rate information comes almost entirely from events below our energy region of interest (97% of the charged current events, and nearly all the elastic scattering and neutral current events, are below our energy window), so in practice these are nearly independent data sets. As we constrain the ^{8}B rate, it cannot be far from the true values, and indeed it is not. In the one-third data sets, the means of the posteriors are all very close to the true value of 4300, and the standard deviation (Subfigure (c)), is around 170, which exactly matches the prior. In the full data sets, the mean and standard deviation of the ^{8}B event rate are a little higher than the true values of $13,400 \pm 510$, around $13,550 \pm 530$. These biases can be neglected as they are small compared to the uncertainty in the rate.

Now let us consider the energy scales and offsets, and the overall rate correction, Figures

7.1 (d)-(i) and 7.2 (d)-(i). First, we observe the energy shifts are typically around -0.03 MeV, which is very close to the true value of 0 MeV, especially considering the individual event energy posteriors are binned by 0.1 MeV. The standard deviation of the energy shifts are around 0.12 MeV, which means the offset is well-constrained by the data. Next if we look at the energy scales, we find these are also close to the expected values of 1.0. The energy scale is negatively correlated with the energy offset (we found correlations between the scale and offset of about -0.10), which explains the slight positive bias of about 2-3%. We also see a small bias in the rate correction, about 2% in the full data set, which is smaller than its prior of 6%. It is a real success that we were able to determine the energy scale to within about 10% (on the full data set), and with a bias of only a few percent, simply by floating the calibration parameters in the fit. The fact that the biases and standard deviations decrease when we increase the statistics from the simulated one-third to full data sets shows the the determination of the energy scale is statistics-limited.

7.3 Sensitivity

As the expected rate of *hep* events is very low, we wonder what limit we could set if in fact there is no *hep* signal. To test this, we repeat the Monte Carlo tests above with rates of 13,400 ^{8}B events and 0 *hep* events. We created 100 data sets, fit each of them, and then found the maximum value of the 95% credible region. These are shown in Figure 7.4. We find that, if the *hep* rate were truly zero, with the full data set we would expect to be able to rule out a rate above approximately 78 events in the data set (this corresponds to a flux of $1.27 \times 10^4/\text{cm}^2/\text{s}$. For comparison, in the Monte Carlo with the expected rate of *hep* events included, the 95% credible region extends up to 135 events in the data set (or a flux of $2.20 \times 10^4/\text{cm}^2/\text{s}$).

7.4 Instrumental cuts

Before we can extract the *hep* signal from the real data, we must remove as many background (non-Cherenkov) events as possible. We first remove many types of backgrounds by applying

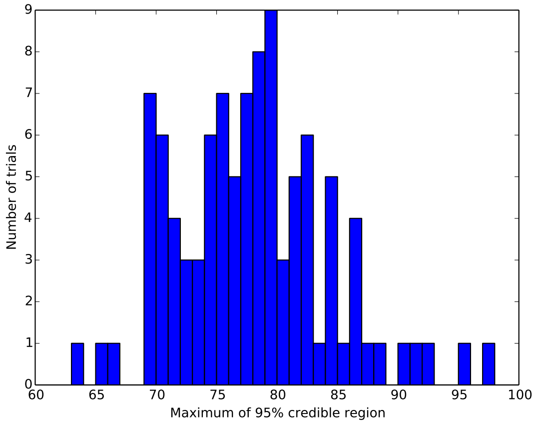

Figure 7.4: Sensitivity of *hep* test. (Maximum values of 95% credible region.)

DAMN cuts (See Chapters 2 and 5) corresponding to the mask `0x0CF56FE1`. These remove instrumental events such as those caused by flashing PMTs, crosstalk in the circuit boards, events in the neck (which may be caused by static discharges near the water-air boundary), and a few others. These also remove muon events and their followers, and events in close time coincidence, which are unlikely to be solar neutrino events.

Next we apply the same sets of cuts used in previous SNO analyses to remove other non-Cherenkov events. There is a cut on isotropy:

$$-0.12 < \beta_{14} < 0.95 \tag{7.2}$$

and in-time ratio:

$$\text{ITR} > 0.55. \tag{7.3}$$

Both of these were used previously in the ^{8}B energy region. Slightly tighter cuts were previously used in the *hep* energy region, but as we have posterior distributions rather than fixed values, using a different cut at different energies would be a bit arbitrary. Additionally, the goal of the tighter cuts was to reduce AVIBs, for which we developed a better cut, so we use the slightly wider cut windows for β_{14} and ITR.

We use the log-likelihood-ratio developed in Chapter 3 to cut AVIBs.

$$\ln(P_{\text{electron}}/P_{\text{AVIB}}) > 100 \tag{7.4}$$

As was also done in previous analyses, we apply a burst cut in addition to the one that appears in the DAMN cuts. We remove any event occurring within 250 ms before or after another event.

We also have a cut on radius because we are not interested in events that occur in the light water. These events could be from radioactivity in the PMTs (at lower energies), elastic scattering neutrino events in the water, so-called "blind flashers", which are PMTs that flash without producing a tell-tale charge in the flashing PMT, neutrons, atmospheric interactions, or perhaps others. Due to fitter resolution, however, we are interested in charged current events that reconstruct slightly outside the AV, so we place the radius cut at

$$r < 625 \text{ cm}. \tag{7.5}$$

We do not have to worry about events in the acrylic itself because radioactive backgrounds are of lower energy, AVIBs are cut with the log-likelihood-ratio, and Cherenkov events that are genuinely inside the acrylic produce relatively little light, so these will also be ignored by this analysis. As will be discussed further in Section 7.7, the efficiency for *hep* events to pass this cut is 99.8%.

7.5 Examination of the real 1/3 data set

Let us take a moment to examine the real data. First, let us look at the posterior distributions for the events. In Figure 7.5, we plot the standard deviation, skewness, and kurtosis of the real events as a function of the average value of their energy posteriors. We see that there is a considerable variety of these values for the events in our data set. This means that the posteriors tend to differ from a Gaussian shape, and may indeed be asymmetric or have large tails. As the events do not all look the same, this vindicates the choice of attempting to incorporate the full energy posteriors of the events, rather than just using the best value or best value and uncertainty.

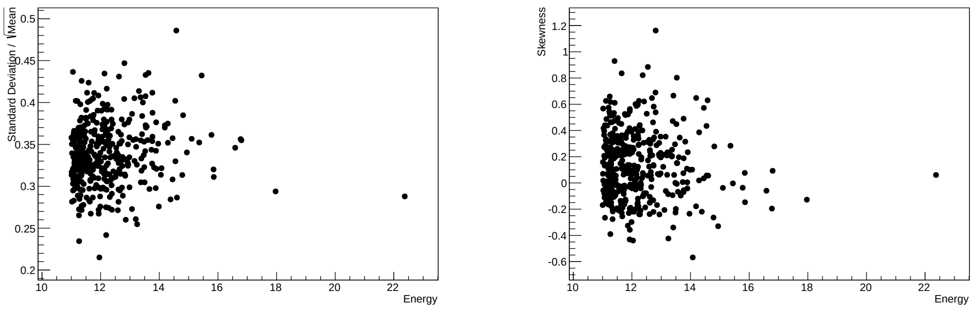

(a) Standard Deviations of the energy posteriors of data events.

(b) Skewness of of the energy posteriors of data events.

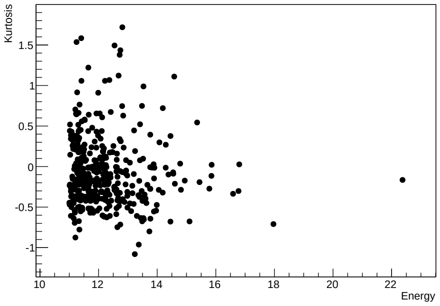

(c) Kurtosis of the energy posteriors o data events.

Figure 7.5: Shape parameters of the energy posteriors of data events.

Cut	Number of events removed by only this cut	Events remaining after cut
DAMN and $N_{\text{PMTs}} > 40$		44819
Burst	1	42979
$r < 625$ cm	616	20833
ITR$>$ 0.55	6	16030
$-0.12 < \beta_{14} < 0.95$	1	15506
AVIB cut	51	8952
11 MeV $< E <$ 23 MeV		$\sim$ 242 (varies by sample)

Table 7.2: Events removed by each cut.

Let us also check how many events are removed by each of our cuts. This is not completely straightforward, again because the energy cutoff is not fixed. Therefore, we will examine events whose mean energy posterior value is above 11 MeV and below 23 MeV (which we determine based on Equation 6.12). A total of 387 events pass this cut. The number of events cut by *only* one cut is shown in Table 7.2. It is unknown exactly how many events were cut by DAMN cuts, as these were removed before even doing fits, however there were over 1×10^8 events initially in the 1/3 data set and roughly 45,000 events of these pass DAMN cuts and a preliminary 40 PMT hit cut.

We show the energies of the events that fail the radius and AVIB cuts in Figure 7.6. We see that the events that fail the radius cut are largely concentrated at lower energy, though there does appear to be a flat component to the distribution. The AVIBs are much more uniform in energy, though there are still slightly more at lower energy.

7.6 Results for 1/3 data set

At last, we run the full analysis on the 1/3 real data set. The posteriors for the rates of *hep*, ^{8}B, and atmospheric events are shown in Figure 7.7. Note that the expected number

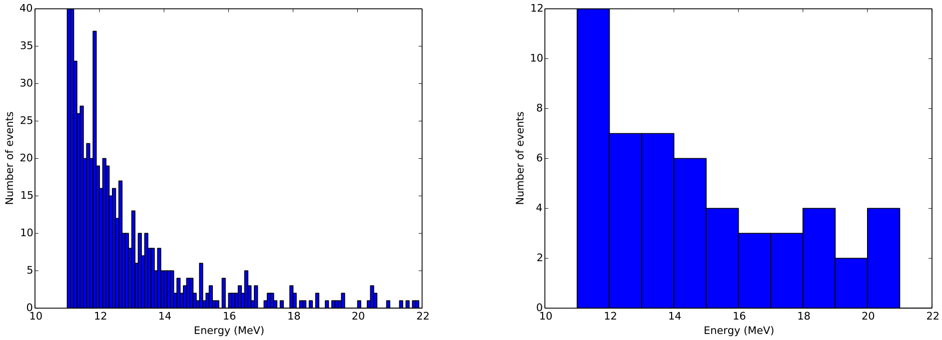

(a) Energy of events that fail the radius cut.

(b) Energy of events that fail the AVIB cut.

Figure 7.6: Energy distributions for cut events.

of *hep* events is 16, while the most probable value in the posterior is about 75 events. This corresponds to a *hep* rate nearly five times larger than the theoretical flux. We find that the 95% credible region excludes 0, and extends from 6 to 191 events in the data set.

As a sanity check (and because we are used to looking at data histograms with PDFs superimposed), we can also combine all the energy posteriors into a single histogram to get an overall sense of the energy distribution. We do this in Figure 7.8. We also include the expected shape of the ^{8}B, *hep*, and atmospheric spectra. To produce these, we first convolved the true PDFs with the uncertainty due to the number of photons produced at a given energy. In order to account for the uncertainty due to the fitter posteriors, we created a model posterior with the same standard deviation, skewness, and kurtosis as the average over all events above. We then doubly-convolved the spectra with this model posterior. From this figure, we can see that the resulting spectra match the data very well, even though this analysis technique is not as involved as our main analysis. Figure 7.9 shows the evolution of the ^{8}B spectrum after each convolution described above.

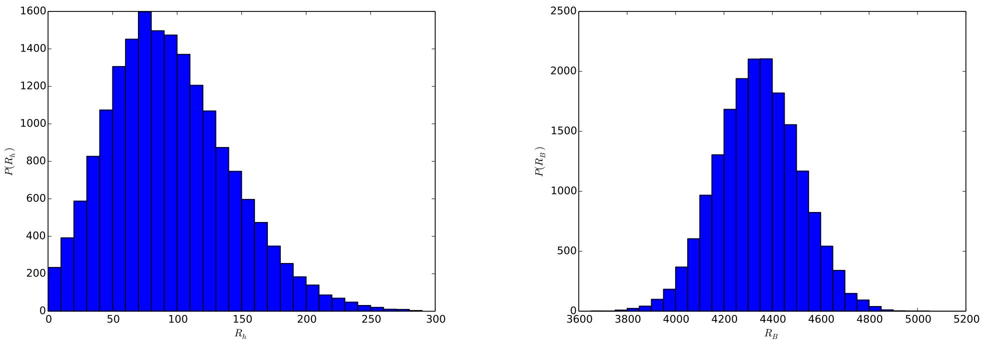

(a) Extracted rate of *hep* events.

(b) Extracted rate of ^{8}B events.

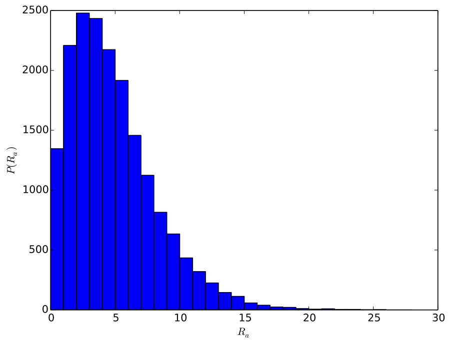

(c) Extracted rate of atmospheric events.

Figure 7.7: Posteriors for the rates of events in real 1/3 data set.

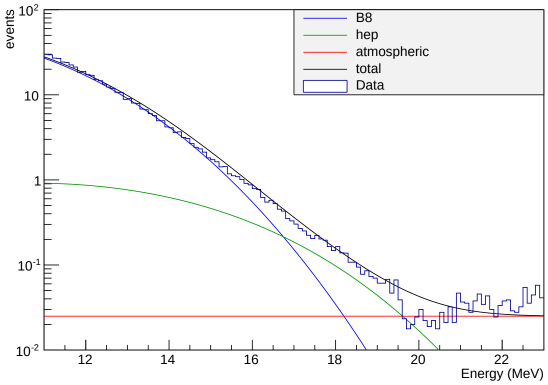

Figure 7.8: Combined histogram of energy posteriors of all events in 1/3 data set.

7.7 *Systematics*

Table 7.3 lists the systematics affecting this analysis.

The MCMC algorithm naturally accounts for uncertainties in the energy scale and offset by allowing these to change from step to step. Therefore, the posterior distribution of the rate of *hep* events includes the uncertainty in these parameters. By placing a prior on the ^{8}B event rate and allowing the atmospheric event rate to float, the uncertainties on these rates are also included by the MCMC.

There are also some uncertainties that affect both the *hep* and ^{8}B rates in a correlated way. These include the live time, fiducial volume sacrifice, and effect of neutrino oscillations on the flux. We incorporate these into the MCMC by allowing a correction to the expected number of events at each step. We constrain this correction to be within the combined uncertainties of these parameters, and in that way these are also included in the *hep* rate posterior.

As discussed previously, the uncertainty on the live time is about 0.3%. The uncertainty on the event rate due to neutrino oscillations is 6%. The survival probability does vary slightly over the *hep* energy region which affects the shape of the spectra slightly, but this is

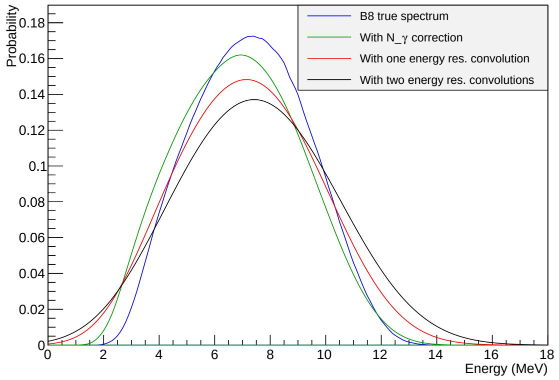

Figure 7.9: The ^{8}B spectrum after each individual convolution to produce Figure 7.8.

too small to affect to be included.

Another uncertainty to consider is that of the fiducial volume. Previous SNO analyses have all needed to do radius cuts in order to avoid backgrounds on the acrylic vessel, but we have removed all such backgrounds with the energy cut and AVIB cut. Therefore, we include *all* the D_2O within the spherical portion of the AV. We may include some additional D_2O from the neck. This would amount to less than a 0.1% increase in the number of events observed. We may, however, also increase our uncertainty by having some events reconstruct outside the 6.25 m cut, or by including addition backgrounds (such as atmospheric events) from the H_2O.

A histogram of the radial position of events from the 1/3 data set with all the instrumental cuts applied (except radius), as well as cutting events with mean energy (estimated using Equation 6.12) below 11 MeV or above 23 MeV, is shown in Figure 7.10. From the figure, we can see that, as expected, the events inside the fiducial volume have a flat distribution in r^3. There is no sign of a bump at the radius of the AV, so this gives us confidence that we are cutting the AV-related backgrounds effectively. We also see a very similar distribution between golden and additional runs, so this is also consistent with our claim that the additional runs contain data of the same quality as golden runs.

Systematic	Uncertainty	Floated in MCMC (Included in posterior)
Energy thresholds	0.1 MeV%	Yes
Energy scale	10%	Yes
Live time	0.3%	Yes
Radius cut sacrifice	0.2%	Yes
Radius cut contamination	2.6%	No
Neutrino oscillations	6.0%	Yes
Atmospheric model	20%	No
AVIB sacrifice	2.0%	Yes
AVIB contamination	0.3 events	No

Table 7.3: Summary of systematics in *hep* signal extraction.

We do see a lot of events at very high radius. As these events pass our energy cuts, their most likely source is flashing PMTs. Usually, a flashing PMT has a distinctive signature caught by the DAMN cuts: one PMT with a lot of charge, at an early time, across the detector from most of the PMT hits. However, for a variety of reasons, a flashing PMT event may be missing the tell-tale high-charge PMT. For example, readout may have been disabled if the PMT had a high rate (it was impossible to turn off voltage to one PMT without powering down many PMTs). Such a flashing-PMT event could reconstruct at a much lower radius than the PMTs for a few reasons. First, if nearly all the light is detected on one side of the detector, then basically the only information about its origin comes from the assumed Cherenkov cone shape, which in this case is invalid. Also, as the fitter is intended for events inside the D_2O, it strongly penalizes light that reaches a PMT from a shallow angle, so moving the event vertex nearer the center often improves the likelihood. It is also possible that low energy events near the PMTs may reconstruct to appear much more energetic, but since we are not interested in events far into the H_2O, we did not look

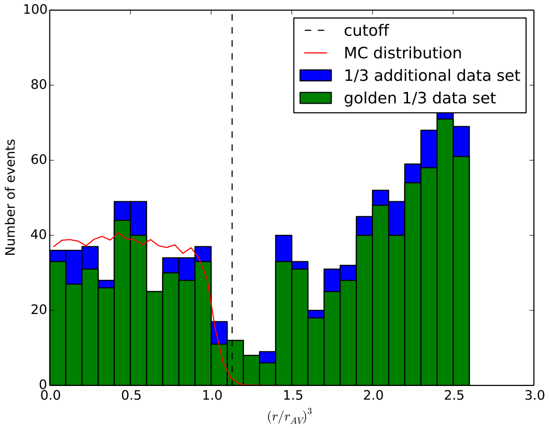

Figure 7.10: Histogram of the best-fit positions of events after all cuts are applied, separated into golden and additional runs. We bin in radius-cubed, in units where the radius of the AV is 1. Also shown is the location of the cutoff (6.25 m), and the expected distribution if the events inside the radius cut are isotropic within the AV.

for this. If this is the case, then some of these high-radius events may in fact be of lower energy than they appear.

Let us refer back to Figure 7.10. In addition to the very-high radius events discussed above, there are 365 events inside a radius of 6.00 m, 21 more inside 6.25 m, and 14 more inside 6.5 m. This means that relative to the the total number of events included, 5.7% reconstruct between the AV and the radial cutoff, and another 3.6% could be included if we pushed the cutoff to 650 cm. In Chapter 4 we performed nearly 100,000 fits to SNOMAN-generated Monte Carlo *hep* charged current events. Of these, 3.1% fit to a position between a radius of 600 and 625 cm, and 0.2% reconstructed between 625 and 650 cm. Therefore, the sacrifice of this cut is only 0.2%. In order to account for 5.7% of the events between 600 and 625 cm, assuming SNOMAN includes all effects that would bias the reconstruction position, we estimate that the additional 2.6% must be from background events in this volume.

Another source of background events we have discussed previously is the possibility of

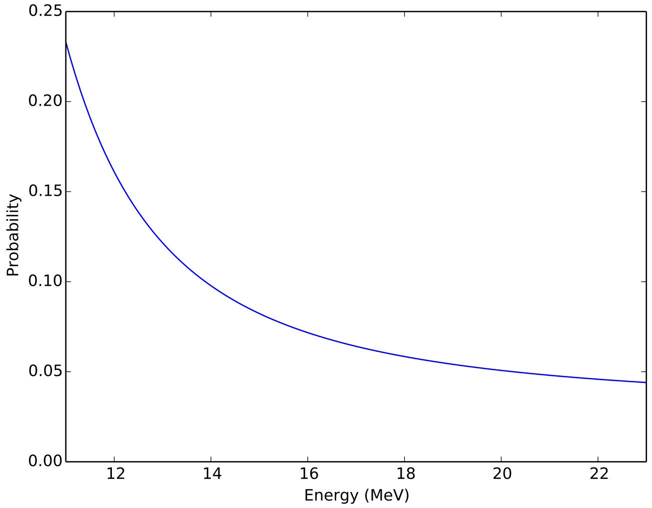

Figure 7.11: Atmospheric event distribution derived from SNOMAN Monte Carlo.

AVIB events passing our AVIB cut. We estimated the number of such events to be 0.3.

A more difficult uncertainty to quantify is that of the shape of the atmospheric background. We used a flat distribution, but it is possible that certain portions (such as $^{12}C^*$ γs) could result in a distribution that increases at lower energy and insidiously mimics the *hep* signal. It is difficult to account for this in a model-independent way. Therefore, we will account for it in a model-dependent way by using an atmospheric spectrum derived from SNOMAN Monte Carlo, and comparing this to the flat atmospheric model. With the model based on SNOMAN (Figure 7.11), we find the distributions shown in Figure 7.12. Here, the rate of *hep* events is peaked around 60 events, which is still considerably higher than the expected rate. However, the most probable value has decreased by 20%. With this atmospheric model, the 95% credible region includes 0 and extends up to 169 events in the data set, which is a 12% decrease. For the purposes of reporting an uncertainty in Table 7.3, we use 20%, but as this is somewhat arbitrary, the more useful characterization of this systematic will be to simply present the posterior distribution with each atmospheric model.

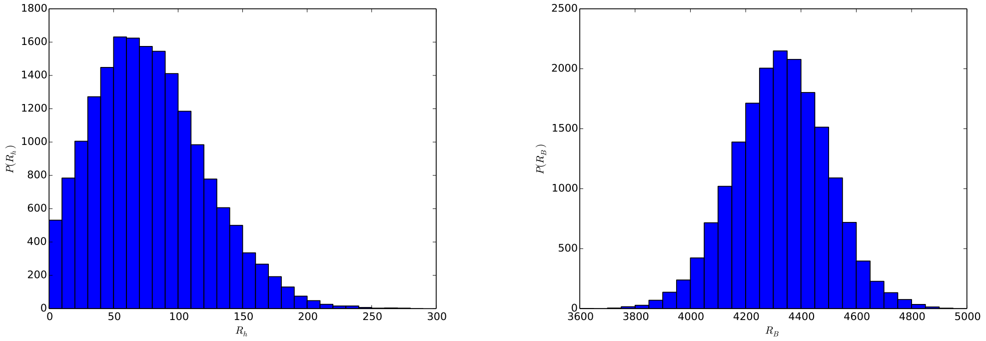

(a) Extracted rate of *hep* events.

(b) Extracted rate of ^{8}B events.

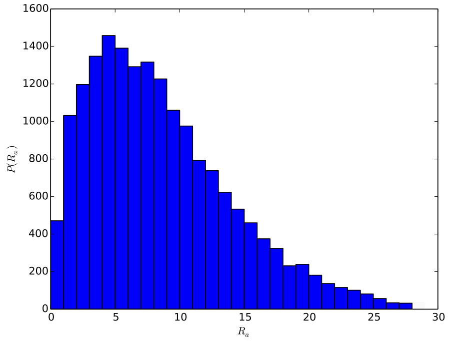

(c) Extracted rate of atmospheric events.

Figure 7.12: Posteriors for the rates of events in real 1/3 data set using the SNOMAN atmospheric distribution.

7.8 Full data results

We proceeded to fit the full data set in the way we described above. Having done so, the ^{8}B rate and overall rate correction both fit more than a standard deviation above their prior constraints. We concluded that the explanation for this behavior was that low-energy background events must be passing the 11 MeV cut, and these look the most like ^{8}B events. These backgrounds are most likely neutron events. We increased the cutoff to 12 MeV,

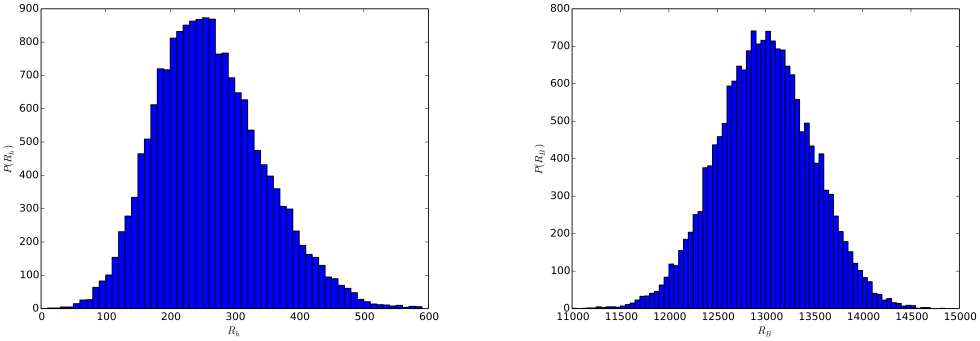

(a) Posterior distribution for *hep* event rate.

(b) Posterior distribution for ^{8}B event rate.

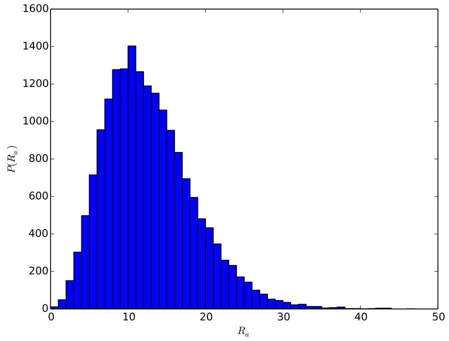

(c) Posterior distribution for atmospheric event rate.

Figure 7.13: Posterior distributions for final signal extraction.

which removed the unexpectedly high ^{8}B rate and high overall rate correction, which we take as confirmation that backgrounds were indeed contaminating the data, and that using a 12 MeV cutoff corrected this.

In Figure 7.13, we show the posterior distributions for the rates of *hep* events, ^{8}B events, and atmospheric events. The expected rate of ^{8}B events is 12900 ± 490, and the expected rate of *hep* events is 49. We find that the ^{8}B rate is precisely as expected, while the *hep* rate is peaked at a value of around 250. This is consistent with the result we saw in the smaller data set (though we now know the one-third data set contained backgrounds at 11 MeV, so one should not overemphasize the comparison).

Before we calculate the final posterior distribution for the *hep* neutrino flux (as opposed to the event rate), we fit the data in two additional ways. First, we fit using the non-flat SNOMAN distribution for the atmospheric background. Second, we fit the data considering only the golden runs and a radial cutoff of 550 cm. The latter fit will give us confidence that our signal is not due to some sort of backgrounds we inadvertently introduced by increasing the fiducial volume or including non-golden runs. Table 7.4 summarizes the energy calibration values, overall rate correction, and 95% credible region for the *hep* flux from these three fits. The fluxes are calculated by reversing the calculation in Section 7.1. We do this for each step in the MCMC output, being careful to adjust for the rate correction at each step. The 95% credible region is calculated by adding the highest-probability bins until 95% of the posterior is included. We can see from the table that the parameters are consistent from fit-to-fit. The uncertainties are larger in the case with the golden data sets and smaller fiducial volume, which is as expected with only 65% the expected number of events. In all fits, the 95% credible region excludes the predicted flux of $7.97 \times 10^3/\text{cm}^2/\text{s}$. Because the golden, small-volume data set is consistent with the full data set, the mismatch with theory cannot be due specifically to backgrounds introduced by including more runs or a larger volume.

Finally, we present posterior distributions for the flux of *hep* neutrinos for each of the three fit conditions in Figure 7.14. We see that the results look quite similar in each case. The choice of atmospheric distribution has some effect on the lower boundary of the 95% credible region, though the peak is about $35 \times 10^3/\text{cm}^2/\text{s}$ in each case. In all cases we rule out a rate of 0 as well as the theoretical rate at greater than 95% probability. The most conservative result, in the sense of being most consistent with expectations, is produced with the SNOMAN atmospheric background. Therefore, we will use this posterior to produce the final result.

As a final check of the fit, Figure 7.15 shows the combined energy posteriors of all events

Fit conditions	Energy offset (MeV)	Energy scale	Rate correction	Φ_h 95% C.R. (10^3/cm^2/s)
Full data, flat atmo.	-0.13 ± 0.28	1.06 ± 0.14	1.010 ± 0.065	17-76
	-0.16 ± 0.14	1.04 ± 0.11		
	-0.14 ± 0.15	1.03 ± 0.12		
Full data, SNOMAN atmo.	-0.17 ± 0.32	1.12 ± 0.12	0.996 ± 0.065	10-72
	-0.28 ± 0.27	1.14 ± 0.14		
	-0.28 ± 0.25	1.12 ± 0.17		
Golden data, flat atmo.	0.14 ± 0.47	1.05 ± 0.18	0.999 ± 0.065	12-112
	0.02 ± 0.49	0.97 ± 0.19		
	0.03 ± 0.36	1.04 ± 0.13		

Table 7.4: Results of signal extraction under three different conditions. The energy calibration parameters are shown for each phase. All uncertainties are the standard deviation of the posterior.

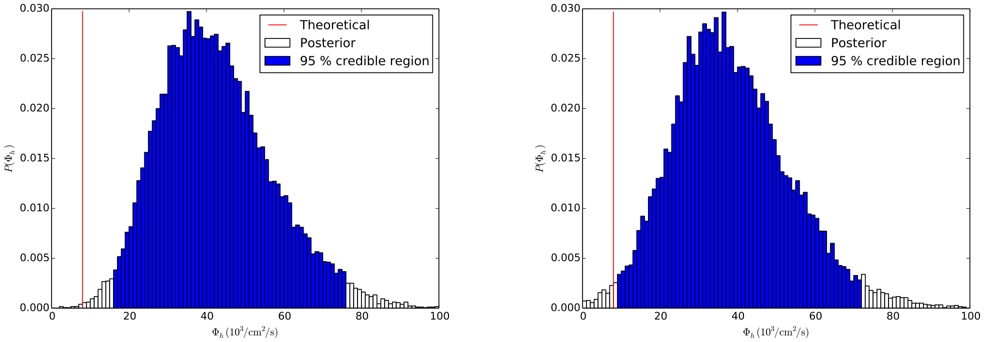

(a) Full data set with flat atmospheric distribution.

(b) Full data set with SNOMAN atmospheric distribution.

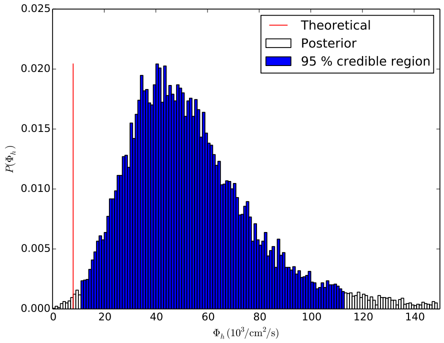

(c) Golden, 550 cm-radius data set, with flat atmospheric distribution.

Figure 7.14: Posterior distributions for the *hep* flux under three different conditions.

along with the expected distributions. As in the 1/3 dataset, we see a reasonably good agreement, even though the technique used to fit the data is very different.

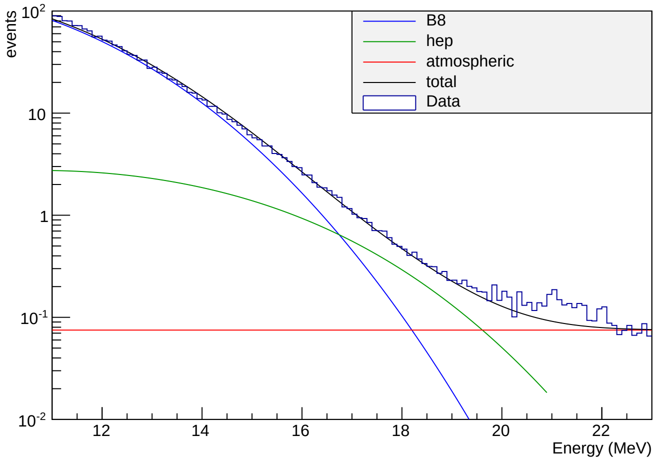

Figure 7.15: Combined histogram of energy posteriors of all events in full data set.

Chapter 8

CONCLUSION

The source of the Sun's energy is the fusion of hydrogen into helium. This process occurs over several possible pathways, each of which produces a total of two neutrinos per ^{4}He produced. Because these neutrinos stream freely from the interior of the Sun without interacting, by measuring these neutrinos we can monitor aspects of the conditions in the interior of the Sun in a way that would be impossible using other methods.

Previous measurements have been made of *pp*, *pep*, ^{7}Be, and ^{8}B neutrinos, but measurements of neutrinos from the rare *hep* interaction,

$$^3\text{He} + p \rightarrow {}^4\text{He} + e^+ + \nu \tag{8.1}$$

had previously been unsuccessful. A measurement of the rate of *hep* neutrinos would complement other solar neutrino measurements because the largest uncertainty on the predicted rates of all other solar neutrinos is the solar composition, while *hep* neutrinos are less dependent on solar composition. Additionally, there is a large theoretical uncertainty on the matrix element of the *hep* reaction, and this could possibly be improved.

The Sudbury Neutrino Observatory (SNO) was a neutrino detector located in Creighton Mine in Sudbury, Ontario that operated from 1999 to 2006. SNO employed 1000 tons of heavy water, D_2O, because two different reactions can occur between a neutrino and a deuteron, one of which detects only electron neutrinos, and one of which equally detects all flavors of neutrinos. SNO demonstrated that only one third of the neutrinos received from the sun are electron flavor, and thereby definitively solved the Solar Neutrino Problem.

SNO was primarily sensitive to ^{8}B solar neutrinos, but as *hep* neutrinos have an endpoint at higher energy than ^{8}B neutrinos, SNO had the potential to detect a *hep* signal. Due to the extremely low flux of *hep* neutrinos, however, previous analyses were unsuccessful at making a detection.

In this study, several methods were employed to improve the sensitivity of a *hep* neutrino measurement. First, we developed an event fitter that uses an MCMC technique to fit position and energy simultaneously, and returns a posterior distribution of energy values. This fitter was demonstrated to reconstruct position with a resolution 10% smaller than the previous best SNO position fitter. This new fitter also has several advantages in energy fitting over the previous best SNO energy fitter. In particular, it performs dramatically better at fitting the energy of events that reconstruct within the AV, which means that much more of the detector's volume becomes usable for analysis. Another feature of the new fitter is it allows a natural way to cut backgrounds such as AVIBs by fitting events to both distributions and comparing the likelihoods in each case. This allows the fiducial volume to include the *entire* D_2O region of the detector, which increases the number of events in the analysis by 30%. Third, the event's posterior distribution contains more information about the event than simply the maximum likelihood value.

Another method we employed was to increase the number of runs used in the analysis. This was accomplished by identifying sets of runs which had previously been excluded but potentially contained good data, such as runs with low-energy backgrounds, runs with calibration sources, runs between the salt and NCD phases, and superficially normal runs excluded for a variety of reasons. We used a modified Wald-Wolfowitz test to compare the distributions of event parameters in the test data sets with the golden data sets, in order to determine whether the distributions of events in the test data sets were consistent with data we know to be good. We also compared the rates of events between the test sets and golden sets to determine if these were consistent. After running these tests on 21 test data sets, we identified an additional 200 days worth of runs consistent with the golden runs, which amounts to an 18% increase in live time.

Finally, we developed a Bayesian analysis that fully used the posterior energy distributions from the event fitter to extract the *hep* signal. Since this analysis has limited statistics, incorporating the full posteriors makes better use of the information we have available. The analysis method also allowed us to naturally incorporate systematic effects like uncertainty in the energy calibration and ^{8}B event rate in a natural way.

We presented the first significant observation of *hep* neutrinos. We measured a rate of *hep* events significantly higher than the theoretical prediction of $7.97 \times 10^3/\text{cm}^2/\text{s}$. We found that the most probable value of the *hep* flux is $34 \times 10^3/\text{cm}^2/\text{s}$, and the 95% credible region extends from $10 - 72 \times 10^3/\text{cm}^2/\text{s}$. A possible cause of this deviation is that the largest uncertainty in the theoretical hep rate comes from the nuclear matrix element. The matrix element calculation involves considerable cancellation of terms, which may account for the mismatch with observation. A reexamination of the theory may provide clues about the origin of the mismatch.

8.1 Improvements for future work

With the benefit of hindsight, there are a few areas where this analysis could be improved. The event fitter could be improved by including a more complicated model of the detector. Specifically, there is some evidence the neck and belly plates affected energy resolution. Also, including even a simple model of the NCDs would surely improve energy biases in the Phase III fits. Finally, a more accurate model of the PMT angular response may help improve the position-dependent energy biases in all three phases. A goodness-of-fit cut could also be developed based on the pulls of the fit parameters. A fit could also be performed with a model of flashing PMTs so these can be cut, as was done with AVIBs.

For the signal extraction, one thing we neglected was the contribution of elastic scattering events, which should be included before publication. While these are a reasonably small portion of the total number of events in our energy window, they will affect the shape of the spectra and overall flux. They may also explain the unexpected presence of events in the 6.0-6.5 m radius region, and could affect the *hep* signal we measured. While elastic scattering would affect the shape of the spectra, if we assume it is only an adjustment to the rate, we can estimate the effect of this on our result. We show in Figure 8.1 the *hep* and ^{8}B rates at

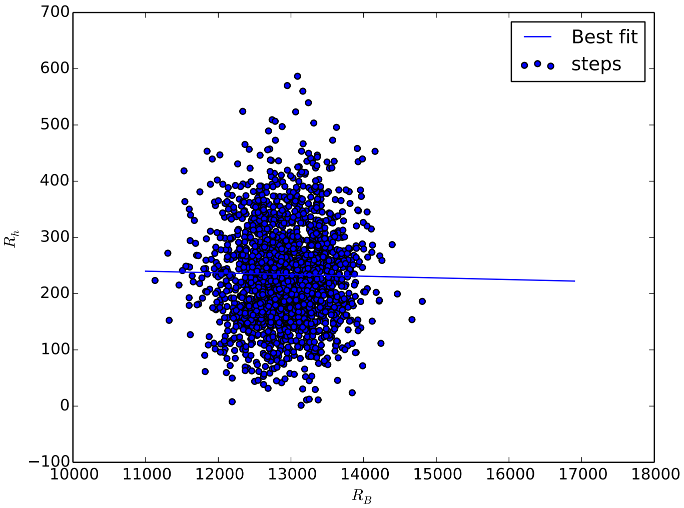

Figure 8.1: Relationship between ^{8}B and *hep* rates in the signal extraction.

every 1000th step in the MCMC, along with a best-fit line. The correlation is only -0.0165, so the *hep* signal barely varies with the rate of ^{8}B events, and we expect a minimal impact on the final result.

8.2 Additional remarks about the fitter

While the event fitter was developed and used as part of a *hep* analysis, its performance was comparable to or better than existing fitters in most ways. Therefore, one could use it for an improved ^{8}B analysis. There are a number of suggestions I would make to someone who was interested in taking that on.

First, I think there are at least two possible ways to go about discriminating neutrons from electrons. One would be to use the relative directions and intensities of the cones, and the fractions in the flat and parabolic parts of the distribution, to try to develop a parameter akin to β_{14} which has different distributions for electrons and neutrons. The second way, and the one I would recommend, would be to come up with a model for the neutrons and, as with the AVIBs, fit all events with both a neutron-like and electron-like distribution, then compare the likelihoods. I'm not sure what the neutron distribution might be, though it could be as simple as comparing a multiple-cone distribution (for a neutron) to a one-cone distribution (for an electron). It also should be mentioned that the three cones share a vertex currently, but in a more complicated neutron model, the gammas might travel a considerable distance before Compton scattering.

To distinguish elastic scattering events from charged current events, a useful parameter is the cone direction. Unfortunately, in its current form the fitter has three cones in three directions. It may be tempting to preferentially use the most intense cone, but due to the symmetry in the fitter, the cones have a tendency to swap occasionally, so this may not be a reliable choice. Perhaps a new parameter could be calculated step-by-step to pick the direction of the best cone, or else the one-cone fit could come to the rescue again.

BIBLIOGRAPHY

[1] K. Abe et al. Solar neutrino results in Super-Kamiokande-III. *Phys. Rev.*, D83:052010, 2011.

[2] B. Aharmim et al. A Search for Neutrinos from the Solar hep Reaction and the Diffuse Supernova Neutrino Background with the Sudbury Neutrino Observatory. *Astrophys.J.*, 653:1545–1551, 2006.

[3] B. Aharmim et al. Determination of the ν_e and total ^{8}B solar neutrino fluxes with the Sudbury neutrino observatory phase I data set. *Phys. Rev.*, C75:045502, 2007.

[4] B. Aharmim et al. Measurement of the Cosmic Ray and Neutrino-Induced Muon Flux at the Sudbury Neutrino Observatory. *Phys. Rev.*, D80:012001, 2009.

[5] B. Aharmim et al. Combined Analysis of all Three Phases of Solar Neutrino Data from the Sudbury Neutrino Observatory. *Phys.Rev.*, C88:025501, 2013.

[6] S. Andringa et al. Current Status and Future Prospects of the SNO+ Experiment. *Adv. High Energy Phys.*, 2016:6194250, 2015.

[7] M. Antonello et al. Measurement of the neutrino velocity with the ICARUS detector at the CNGS beam. *Phys. Lett.*, B713:17–22, 2012.

[8] Martin Asplund, Nicolas Grevesse, A. Jacques Sauval, and Pat Scott. The chemical composition of the Sun. *Ann.Rev.Astron.Astrophys.*, 47:481–522, 2009.

[9] Nigel R. Badnell, M.A. Bautista, K. Butler, F. Delahaye, C. Mendoza, et al. Up-dated opacities from the Opacity Project. *Mon.Not.Roy.Astron.Soc.*, 360:458–464, 2005.

[10] John N. Bahcall. *Neutrino Astrophysics*. Cambridge, UK: Univ. Pr. (1989) 567p, 1989.

[11] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. New solar opacities, abundances, helioseismology, and neutrino fluxes. *Astrophys. J.*, 621:L85–L88, 2005.

[12] John N. Bahcall, Aldo M. Serenelli, and Sarbani Basu. 10,000 standard solar models: a Monte Carlo simulation. *Astrophys.J.Suppl.*, 165:400–431, 2006.

[13] G. Bellini et al. Neutrinos from the primary protonproton fusion process in the Sun. *Nature*, 512(7515):383–386, 2014.

[14] H. Bethe and R. Peierls. The 'neutrino'. *Nature*, 133:532, 1934.

[15] H.A. Bethe. Energy production in stars. *Phys.Rev.*, 55:434–456, 1939.

[16] H.A. Bethe. Moliere's theory of multiple scattering. *Phys.Rev.*, 89:1256–1266, 1953.

[17] Xin Chen. *Monte Carlo Simulations and Analysis of Backgrounds in the Sudbury Neutrino Observatory*. PhD thesis, Oxford University, 1997.

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

[19] Monica Dunford. *Measurement of the 8B Solar Neutrino Spectrum at the Sudbury Neutrino Observatory*. PhD thesis, University of Pennsylvania, 2006.

[20] H. Ejiri. Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei. *Phys. Rev.*, C48:1442–1444, 1993.

[21] Friedman, J.H. and Rafsky, L.C. Multivariate Generalizations of the Wald-Wolfowitz and Smirnov Two-Sample Tests. *Ann.Stat.*, 1979.

[22] Walter J. Hamer. Theoretical Mean Activity Coefficients Of Strong Electrolytes in Aqueous Solutions from 0 to 100 degC. Technical report, Institute for Basic Standards, National Burueau of Standards, 1968.

[23] Chris William Howard. *A search for hep neutrinos with the Sudbury Neutrino Observatory Detector*. PhD thesis, University of Alberta, 2010.

[24] J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing In Science & Engineering*, 9(3):90–95, 2007.

[25] J.V. Jelley. *Cerenkov Radiation and its Applications*. Pergamon Press, 1958.

[26] O.S. Kirsebom, S. Hyldegaard, M. Alcorta, M.J.G. Borge, J. Buscher, et al. Precise and accurate determination of the B-8 decay spectrum. *Phys.Rev.*, C83:065802, 2011.

[27] Josh Klein and Mark Neubaer. Using Time and Angle Information in Reconstruction and the SNOMAN Path Fitter. Technical report, University of Pennsylvania, http://www.sno.phy.queensu.ca/sno/str/SNO-STR-2000-022.ps.gz, 2001.

[28] Joshua R. Klein. Tuning the High Level Cuts for LETA v.3.0. Technical report, University of Pennsylvania, 2006.

[29] Ryan Francis MacLellan. *Energy calibration for the solar neutrino analysis of all three phases of the Subdury Neutrino Observatory.* PhD thesis, Queen's University, 2009.

[30] T.S. Park, L.E. Marcucci, R. Schiavilla, M. Viviani, A. Kievsky, et al. Parameter free effective field theory calculation for the solar proton fusion and hep processes. *Phys.Rev.*, C67:055206, 2003.

[31] Wolfgang Pauli. Pauli's letter of the 4th of december 1930, December 1930.

[32] Carlos Pena-Garay and Aldo Serenelli. Solar neutrinos and the solar composition problem. 2008.

[33] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov. Phys. JETP*, 26:984–988, 1968. [Zh. Eksp. Teor. Fiz.53,1717(1967)].

[34] Andrew Renshaw. Solar Neutrino Results from Super-Kamiokande. *Phys. Procedia*, 61:345–354, 2015.

[35] SNO 3-Phase Working Group. 3-Phase Unidoc. Technical report, SNO Collaboration, 2011.

[36] SNO Collaboration. *SNOMAN companion.* `http://www.hep.upenn.edu/~geneb/web_snoman/doc/html/companion_frames.html`.

[37] SNO Collaboration. The SNOMAN User's Manual Version 5.03. Technical report, SNO Collaboration, 2006.

[38] SNO hep Analysis Group. The hep Analysis: Unified Document. Technical report, SNO Collaboration, 2006.

[39] Susan G Stanton, R Pecora, and Bruce S Hudson. Resonance enhanced dynamic rayleigh scattering. *The Journal of Chemical Physics*, 75(12Renshaw:2014awa):5615–5626, 1981.

[40] W. T. Winter, S. J. Freedman, K. E. Rehm, and J. P. Schiffer. The B-8 neutrino spectrum. *Phys. Rev.*, C73:025503, 2006.

[41] L. Wolfenstein. Neutrino Oscillations in Matter. *Phys.Rev.*, D17:2369–2374, 1978.

Appendix A

SNO+ SLOW CONTROLS

As I described briefly in chapter 2, the SNO detector is not retired, but rather is currently being revamped for a new experiment dubbed SNO+. The most important change between SNO and SNO+ is that where the acrylic vessel was filled with D_2O for SNO, it will be filled with a liquid scintillator for SNO+. This seemingly straightforward change allows the same detector to be used for an array of new science capabilities, but also required a number of significant upgrades to many parts of the detector.

One required upgrade is due to mechanics. Heavy water is literally heavier than normal water, and therefore ropes were required to hold the acrylic vessel *up* under this weight. Meanwhile, the liquid scintillator chosen for SNO is mostly composed of oil and is therefore less dense than water, so for SNO+ a new set of ropes was installed to hold the acrylic vessel *down*.

Liquid scintillator (in the case of SNO, a solution in linear alkyl benzene of some fluorescent chemicals and, depending on the phase of the experiment, additional solutes such as tellurium and surfactant) produces more photons per MeV of kinetic energy when a charged particle passes through it. This lowers the threshold for charged particles to be detected, which broadens the scope of the detector to seeing perhaps *pp*, *pep*, ^{7}Be, and *CNO* neutrinos from the sun. Additionally, *anti*-neutrinos from reactors and the interior of the earth will now be accessible. Finally, and perhaps most convincingly for funding agencies, by loading a double-beta decay isotope such as ^{130}Te into the liquid scintillator, the detector has the capability to search for neutrinoless double-beta decay.

A consequence of this increased sensitivity to lower energies (and the corresponding increase in the rate of events *and* number of photons per event) is that the readout rate of the electronics needed to upgraded considerably, which has been a major undertaking.

Meanwhile, plants to purify the liquid scintillator, tellurium, and pure water have all needed to be built.

In the effort to upgrade these various aspects of the detector, the University of Washington group has been in charge of a system called the slow control.

A.1 Slow Control

The slow control system is so named because, while physics events take place on time scales of nanoseconds, we wish to monitor and control certain aspects of the detector on time scales of seconds or minutes. Such slowly varying quantities include temperatures, voltages, warning signals, rope tensions, and water pressure.

Most of these readings are recorded by four Acromag I/O servers running a Linux Gentoo operating system and containing IOS-320 analog input cards, plus one IOS-408 card used for digital output. Other readings are imported from an assortment of other sources including a DeltaV system which records the acrylic vessel hold-up and hold-down rope tensions. Additional systems measure the water depth, the acrylic vessel position, the temperature on deck and of the water, the humidity, and the status of the cover gas system.

In more detail, 11 racks of electronics responsible for the PMTs and one for timing and monitoring each have five power supplies, whose voltages are recorded directly by the I/O servers. Additionally, 19 crates of electronics are distributed among these racks and each has three current readings that are monitored in the same way. These readings allow instant feedback if a power supply fails.

Each of the 19 crates contains a card called an XL3 that, among other things, reports signals indicating that the voltage of the PMTs is okay and that the temperature of the crate is okay. The slow control system monitors these signals and also has the capability to remotely reset the XL3s in case of error.

Embedded in the rock around the detector are 14 coils used to compensate for the Earth's magnetic field. The slow control system can turn these coils on and off remotely, and also monitors signals from the coils that indicate whether they are on and whether there is some sort of failure.

The I/O servers also monitor whether the emergency stop has been pressed, and whether the mine power has failed and the lab has engaged the Uninterruptible Power Supply.

The I/O servers collect data at intervals of roughly 5 seconds, and store the data locally as documents in a CouchDB database, which was chosen for its flexibility and scalability. The I/O servers also compare each reading to a set of thresholds (also stored in a database) and issue alarms if necessary. In addition to being stored in a database, the alarms are sent by email to a list of interested parties, and are displayed by a web interface detailed below. All data are replicated to a central database and a backup database which are then accessible by users. Data is also compiled into 1-minute and 15-minute summaries for ease of access and analysis.

A web interface allows detector operators to view data from the central database and change alarm thresholds or disable alarms in real time. It also summarizes alarms and issues audio alarms when appropriate. Another feature of the web interface is the ability to plot recent trends in the data.

The "control" part of slow control is that the compensation coils can be turned on or off remotely, and the XL3 cards can be reset remotely. Relays ensure that in case of a loss of network or power, the states of these circuits remain unchanged.

Appendix B

SCINTILLATOR SCATTERING MEASUREMENTS

A considerable amount of my time in grad school was spent in service to SNO+ working on measurements of light scattering in various liquid scintillator cocktails. The goal was to parametrize the scattering as being a combination of Rayleigh (or possibly Mie) scattering and absorption/reemission. As these types of scattering have different dependence on wavelength, angle, and polarization, we scanned through these values for each sample.

B.1 Data collection

The apparatus used to measure scattered light as a function of wavelength, scattering angle, and input and output polarization is shown in figure B.1. It is housed in a dedicated dark box. A polished, cast acrylic cylinder (1.5 inch OD, 1.25 inch ID) acts as the sample holder. An acrylic stopper is used to cover the sample to prevent contamination from entering the sample. The sample holder is mounted on a 0.5 inch diameter post that acts as the pivot of a turntable.

On one end of the turntable is mounted one of a set of interchangeable lasers. The laser beam is aligned to pass perpendicular to the axis of the sample holder. On the other end of the turntable is mounted a PMT with ND filters totaling a factor of about 10^8 attenuation of the laser light. This “normalization” PMT is used to measure the total light that passed through the sample.

Beside the turntable, there is mounted a stationary PMT with a baffle designed to accept only light that was scattered near the axis of the sample holder. We call this the “scattering” PMT. In front of the baffle there is a shutter to protect the PMT when the dark box is opened, so that the PMTs can remain under high voltage. In front of the shutter is an adjustable

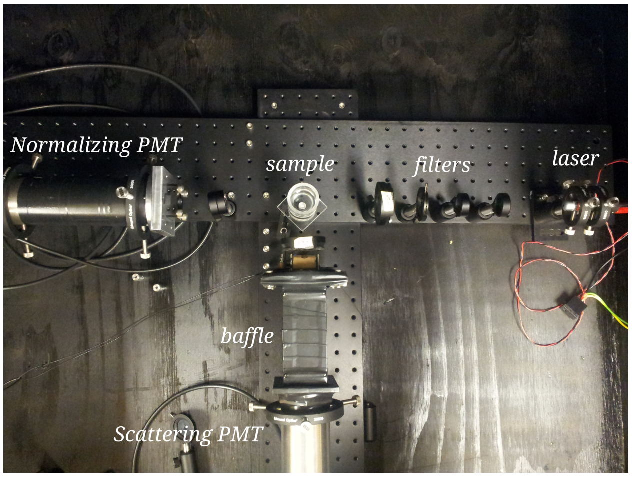

Figure B.1: A photo of the scattering apparatus. On the right is the laser, incident on several filters, the sample in the center, and finally the normalization PMT on the left. At the bottom is the scattering PMT with the baffle above it.

polarizing filter to select only vertical or horizontal polarizations of scattered light.

The baffle is constructed out of black acrylic, a material chosen to minimize diffuse reflections (since specular reflections are minimized by the baffle design). As is shown in figure B.2, the baffle contains six apertures evenly spaced between a shutter and the PMT. The inside apertures were cut with a 45° bevel so that reflections off the cut surfaces cannot reach the PMT. The first and last apertures are cut at a steeper angle (60°) due to machining limitations on the thicker acrylic used. The directions of the first and last bevels are important to prevent accidentally concentrating rather than obscuring unwanted light. The orientation of the bevels on the inner vanes is not critical, but was chosen to evenly space the actual apertures. The dimensions of the baffles are such that light passing through an aperture cannot reflect off a wall and pass directly through the next aperture.

The turntable also features space to insert various components between the laser and the sample holder or between the sample holder and the “normalization” PMT. A typical set of components between the laser and sample holder would be a wavelength filter (to remove

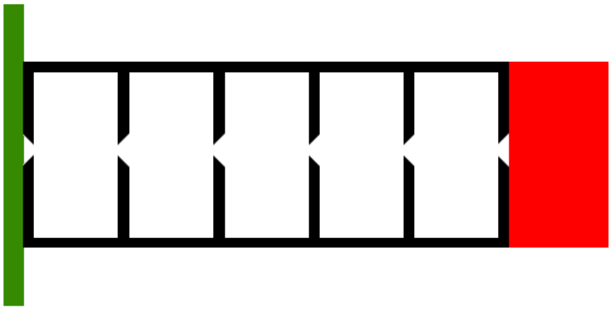

Figure B.2: Design of the baffle. The baffle (black) contains apertures in six vanes which accept light only along their axis. A shutter (green) prevents light from entering the baffle except during measurements. The scattering PMT (red) is flush with the baffle.

unwanted harmonics from frequency-doubled lasers), a 45-degree polarizing filter (to match the vertically-polarized incident intensity to the horizontal), ND filters to adjust the total light intensity, an iris to constrict the beam size (and intensity), and an adjustable polarizer to select vertical or horizontal incident polarizations of light.

Occasionally, it is necessary to use an additional ND filter between the sample holder and the normalization PMT. This is typically required for 650 nm measurements, where the existing ND filters in front of the normalization PMT have much higher transmissivities than for other wavelengths. Meanwhile, less light is scattered at 650 nm, which increases the mismatch between PMT count rates. Conversely, in some cases it was necessary to place an ND filter in front of the scattering PMT. This was typically required only at 405 or 447 nm, where scattering is highest and the ND filters in front of the normalization PMT have lowest transmissivity.

A typical measurement proceeds as follows. The sample is placed into the sample holder and given a day to let any bubbles or possible debris float or sink out of the scattering volume. A laser is mounted onto the turntable and aligned with the center of the sample holder, at the correct height for the baffle. The beam intensity is adjusted by controlling the laser's current, adjusting the iris, or adding filters to the beam line. The optimum laser intensity is chosen so that under the highest scattering conditions (45-degree scattering angle, vertical incident and scattered polarization), the PMTs measure about 100-300 kcounts/min. This is a sufficiently high count rate to allow precise measurements under subsequent conditions, yet small enough to avoid problems due to dead-time.

Once the laser is set up, measurements are taken at four scattering angles (45°, 60°, 75°, and 90°) and four polarizations (each combination of vertical and horizontal incident and scattered) and a "dark" rate measurement at each angle. The dark measurement involves putting an absorbing beam block in the beam path after as many optical components as possible but before the sample. In this way, some of the light that happens to scatter off the walls of the dark box and make it past the baffle can be subtracted in addition to the internal dark rate of the PMTs.

The process is then repeated for the other lasers (wavelengths 405, 447, 473, 532, and 650 nm).

B.2 Analysis

We will now discuss the various corrections made to the data to convert our measurements into absolute quantities that do not depend on our apparatus.

In this apparatus, an ADC is used to count PMT pulses. The ADC has a dead time of 7.5μs after each count. At the count rates employed (maximum 400k/min total for both PMTs), this never amounts to more than 5% change in the measured rates, and a smaller difference in ratio of scattered counts to normalization counts. Nevertheless, it is possible to correct counts based on a known dead time using this equation, which assumes counts are randomly distributed in time:

$$r = \frac{m}{1 - md,} \tag{B.1}$$

where r is the true rate, m is the measured rate, and d is the dead time. This correction is made to all the measurements.

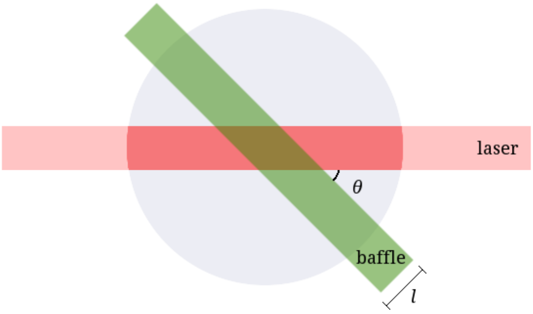

Figure B.3: Top view cartoon of scattering in the sample (grey), in which refraction is ignored. The laser beam (red) defines a cylindrical volume from which light scatters. The baffle accepts light from a clindrical volume (green). The intersection of these volumes is proportional to $\sin^{-1}\theta$ where θ is the angle between the laser beam and baffle acceptance region. This is the source of the geometric correction.

The appropriate dark rates are subtracted from each measurement, and scattered counts are divided by normalization counts to remove dependence on absolute laser brightness. Also, all filters (including the scattering polarizer) have wavelength-dependent transmissivities published by the manufacturer. The fraction of light scattered is corrected for these transmissivities. Filters between the laser and sample affect the light reaching the two PMTs equally, so they are ignored.

A geometric correction is required. The source of this correction is that at angles other than 90°, the scattering PMT sees a larger portion of the laser beam. As shown in figure B.3, if the laser beam and the region the baffle accepts scattered light from both have a uniform (but arbitrary) cross-section, the intersection of these two regions has a volume proportional to $\sin(\theta)^{-1}$, where θ is the scattering angle. Therefore, any type of scattering will be scaled by this factor, and we must therefore multiply all measurements by $\sin(\theta)$ to remove this dependence.

The amount of light collected by the scattering PMT depends on the solid angle between

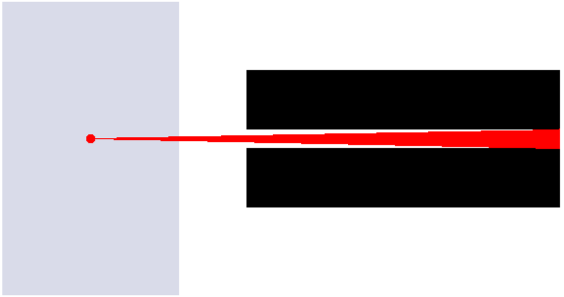

Figure B.4: The amount of light collected by the scattering PMT is proportional to the solid angle (red triangle) between the center of the scattering volume (red spot) and the baffle (black) aperture closest to the PMT. One can also see how Snell's Law changes the solid angle as scattered light leaves the sample.

the center of the sample and the baffle aperture closest to the PMT. See figure B.4. The area of the baffle apertures is 0.2445 cm^2 and the distance between the center of the scintillator sample and the scattering PMT is 24.13 cm. The solid angle is therefore 0.00042 sr. Because the angle is small, the solid angle from points slightly away from the center of the sample will be negligibly different.

The amount of scattered light measured also scales with the path length of the laser beam through the baffle acceptance region (BAR), l in figure B.3. The width of the BAR was measured by putting an LED on a micrometer stage and scanning horizontally. Figure B.5 shows the amount of light collected as a function of lateral position. The large dynamic range (the order of 10^7) was accommodated by setting the LED brightness such that that the PMT count rate was appropriate for some small range. Measurements were taken in eight such ranges with sufficient overlap (3 data points on either end) to determine the scaling between consecutive LED brightnesses.

We define the width of the BAR as the FWHM, which makes sense to use because the finite width of the LED convolutes the real BAR into the measured BAR. Assuming the LED

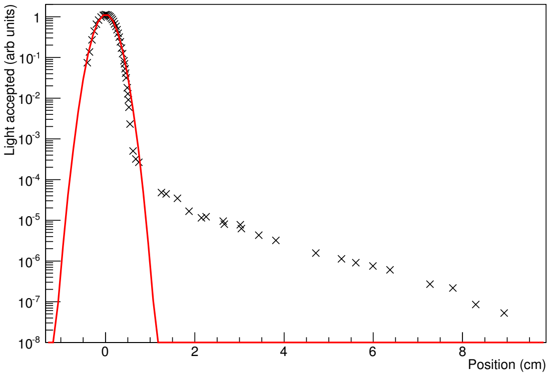

Figure B.5: A plot of the amount of light collected by the scattering PMT as an LED on a micrometer stage is scanned horizontally at the distance of the center of the sample.

is narrower than the BAR (which it is), the FWHM should be unaffected by this convolution. The FWHM of the BAR is measured to be 0.436 cm.

We can also predict the FWHM of the BAR from the geometry of the baffle. The baffle is 5 in long with apertures 0.1 in wide, and the LED is 4.5 in in front of the baffle. Based on these dimensions, the FWHM of a pointlike LED is expected to be 0.483 cm, and will decrease by 0.02 cm for each cm the LED is closer to the baffle. Based on the measured value and the predicted value, we will use 0.46 ± 0.03 cm.

Adjustments must also be made for the difference in index of refraction between the sample and air (the acrylic is assumed to have the same index as the sample). In fact, there are two separate corrections because of the indices of refraction. The first correction is because the sample is in a cylinder rather than a sphere, so light scattered in any direction away from the horizontal is refracted in a vertical plane according to Snell's Law. This can be seen in figure B.4. Therefore, light accepted by the baffle actually must have been scattered within a smaller solid angle than reported above, by a factor of the index of refraction.

The second correction due to the index of refraction is due to distortion in the horizontal plane. Naively, it would seem that light scattered in the horizontal plane should not be

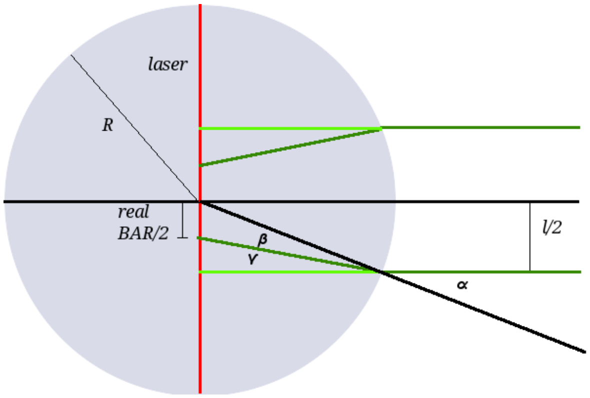

Figure B.6: A top view of the sample (grey) exaggerating how the Baffle Acceptance Region l is distorted by Snell's Law. The baffle only collects light scattered from between the dark green lines, rather than the light green lines. The laser beam is assumed to be infinitely narrow.

refracted because it is exiting normal to the surface. However, the finite BAR allows some of the light to intersect the surface at an angle smaller than 90 degrees. See figure B.6. Consider a photon that travels on the dark green path. Because it started slightly away from the center of the sample (grey), it does not approach the boundary normal to the surface, but at an angle β. It then exits at angle α, related to β by Snell's Law. The photon's path was chosen to just skim the edge of the baffle, so it runs parallel to the center of the BAR a distance of $l/2$, where l is the measured BAR. The photon's starting position, then, represents the furthest position at which scattered light is actually collected, or the "real" BAR. All angles are small (less than 0.01 rad), so the various quantities are related in the following way:

$$
\begin{aligned}
\alpha &= \frac{l}{2R} \\
\alpha &= n\beta \\
\alpha &= \beta + \gamma \\
real\,BAR &= l - 2R\gamma.
\end{aligned}
\tag{B.2}
$$

The first line is the small-angle approximation of the sine function. The second line is the small-angle approximation of Snell's Law. The third and fourth line come from geometry, with the small angles again invoked in the fourth line. Together, these give the relationship $real\,BAR = \frac{l}{n}$.

For completeness, the laser beam is refracted in a way similar to the BAR when it enters the sample. However, this does not have a considerable effect on the amount of light the scattering PMT measures. This refraction simply squeezes the beam into a slightly smaller cross-sectional area, but doesn't cause any of the beam to be moved into or out of the BAR.

Applying all corrections gives the following equation:

$$
corrected\,scattering = \left(\frac{DT(scatter) - DT(dark_s)}{DT(norm) - DT(dark_n)}\right)\sin\theta\left(\frac{FT_n}{FT_s}\right)\frac{n^2}{l}\frac{1}{d\Omega'} \tag{B.3}
$$

where DT is the dead time correction, $scatter$ and $norm$ are the two PMT count rates, $dark_s$ and $dark_n$ are the dark count rates for each PMT, $\sin\theta$ is the geometric correction, FT_n and FT_s are the filter transmissions of all filters in front of the normalization and scattering PMTs, respectively, n is the index of refraction of the sample, l is the measured BAR width, and $d\Omega'$ is the solid angle of scattered light detected. The resulting quantity has units cm^{-1}sr^{-1} and can be interpreted as the fraction of light scattered in the given direction for each cm of path length.

The data are then fit to the following equations:

$$\begin{aligned} VV &= \frac{R}{\lambda^4} + A(\lambda) \\ VH &= A(\lambda) \\ HH &= \frac{\cos^2\theta R}{\lambda^4} + A(\lambda) \\ HV &= A(\lambda) \end{aligned} \tag{B.4}$$

The "R" terms in the first and third equation are Rayleigh scattering, which has a λ^{-4} dependence. Rayleigh scattering also has a polarization-dependent angular dependence, such that for vertical in, vertical out ("VV"), there is equal scattering in all directions, for horizontal in, horizontal out, ("HH") there is a $\cos^2\theta$ distribution (where θ is the angle between incident and scattered light), and for cross-polarizations ("VH" and "HV") there is no scattering.

The "$A(\lambda)$" terms in equations B.4 are absorption-reemission, which in our model is isotropic and independent of polarization, but does have an unknown wavelength dependence.

Once we have determined optimized parameters for Rayleigh scattering (x1) and absorption-reemission (x5), it is necessary to convert these numbers to a useful form for comparison. A common way to report scattering is as a scattering length. To calculate scattering length from the parameters, one must integrate over the solid angle, average over incident polarization, sum over scattered polarization, and finally invert. (One sums over scattered polarizations because a photon could scatter with either polarization and we want to include all scattering. One averages over incident polarizations because only half of the incident photons will have each polarization.) The integrals give a factor of 8π for absorption-reemission (because of the two scattered polarizations) and $\frac{16\pi}{3}$ for Rayleigh scattering (because of the angular dependence).

$$scattering\,length = \frac{1}{\int (scattering\,fraction) d\Omega} \tag{B.5}$$

B.3 Error budget

There are several possible sources of error. The statistical error is measured to be about 1.5% over consecutive repeated measurements, independent of count rate.

Another source of error is reflections off the surfaces of the sample holder. The intensity of such reflections can be calculated analytically from this equation:

$$R = \left(\frac{n' - n}{n' + n}\right)^2, \tag{B.6}$$

which assumes the light is normally incident on the interface. Using $n_{air} = 1$ and $n_{sample} = n_{acrylic} = 1.5$, 4% of the light is reflected back into the sample when exiting, and therefore gets a second chance to scatter. This 4% also does not reach the normalization PMT, so the total correction to the scattered intensity is 8%. Reflections from other interfaces are either negligibly small or affect both the scattering and normalization measurements equally.

Another source of error is variations in transmissivity of the filters over their surfaces (due, for instance, to variations in thickness). This is tested by moving filters slightly and comparing values of the intensity. The variation from doing this was not noticeably larger than statistical noise (about 1.9% total).

Similarly, the sample holder is not perfect, so there might be variations in the amount of light scattered from the surface from different angles of incident light. To check this, the sample holder was rotated while leaving the laser fixed, and variations in scattered light were measured. The variations were found to be similar to statistical noise for scintillator samples, about 2.4%.

The wavelength of three of the lasers are quoted within +/-5 nm. Using an Ist-Rees Laser Spectrum Analyzer, the wavelengths were found to actually be 402nm, 447nm, 473nm, 533nm, and 661nm. The 402nm laser had a spectrum approximately 2 nm wide (full width), but the other four were less than 1nm wide. These small deviations from the nominal wavlengths do not make a considerable difference even for Rayleigh scattering which scales with λ^{-4}, but the filter transmissivities can vary significantly over this range, particularly near 405nm. The absorptive ND4 filter in front of the normalization PMT has about a 20% lower transmissivity at 402nm than 405nm, which results in a 20% increase in scattered light after correction.

The gain of the PMTs can drift over time due to temperature fluctuations or other causes. According to the PMT data sheet, the gain varies by about 0.2%/K. The data sheet also claims the sensitivity varies about about 1% after a change of count rate. The temperature in the lab varied by perhaps 5°C over the course of the measurements, so this introduces a 1% error.

Differences between the scattering and normalization PMT sensitivities could artificially inflate or suppress the measured absolute scattering. To check this, immediately after the tellurium scintillator was measured, the PMTs were ramped down and swapped. They were ramped up again, and after a weekend the set of scattering measurements was repeated. They differed by about 13%, meaning the PMTs have gains that differ by about 7%.

Since we are dealing with fluorescing agents, one might be concerned about changes in wavelength of emitted light affecting the amount of light measured. Absorption-reemission is measured to be highest by far for the 405nm laser, so this is the wavelength where a change in wavelength of scattered light is most possible. According to the PMT data sheet, the sensitivity is maximum around 420nm, and is about 10% lower at about 370nm and 470nm. Therefore, even a 50nm increase in wavelength would make less than a 5% change in PMT sensitivity at this wavelength.

The ND filters can have a larger change in transmissivity, but most of the ND filters are in front of the normalization PMT, which measures light that did not scatter (and therefore did not change wavelength), so this will not affect our measurements. In some measurements (for samples containing bis-MSB or nano-particles), it was necessary to place an ND2 in front of the scattering PMT. The transmission of the ND2 filter is actually 40% higher at 425nm than 405nm, and 60% higher at 445nm than 405nm, so this change is considerable.

Another source of uncertainty is the exact orientation of the polarizing filters. The polarizers are plastic films sandwiched between glass plates, with the rough location of the polarization direction marked on an edge. To calibrate to the precise polarization direction, the two filters used in these measurements were mounted into rotation stages. An LED was incident on the two polarizers in succession, and finally a PMT. One of the rotation stages was rotated through 180°, measuring the transmitted light every 10°. The resulting data were fit to a $\cos^2\theta$ curve to find at what angle they were precisely aligned. One of the polarizers, in its rotation stage, was then flipped front-to-back, and then this process was repeated. These two offsets give the absolute location of the polarization axis of each polarizer. The uncertainty in this process should be less than 1°, but even if it were off by 5°, the error in transmission would be less than 1%.

There is also some uncertainty in the exact angle of the turntable. It is probably less than 2°, with less than 1° difference measurement-to-measurement. At 45°, an error of 3° would result in a change in the geometric correction of about 5%, which would affect all measurements at this angle. Additionally, the Rayleigh scattering for HH scattering would be off by 10% (in the opposite direction, so actually only 5% after both errors). At other angles, the errors introduced to HH would be a larger fraction, since HH should vanish at 90° scattering angle.

Another source of error is stray reflections off of the various optical components and the walls of the dark box, or even light from outside the box that might reach a PMT. In measurements of the scintillator samples, these reflections were negligible.

B.4 Discussion

We found that the model of Rayleigh scattering plus absorption-reemission worked well in some cases but not so well in others. For example, see Figure B.7, where the model fit well. It also turns out to be known [39] that the magnitude of Rayleigh scattering is enhanced by absorption. This is easy enough to include in the model by not forcing the λ^{-4} dependence (we can still tell absorption-reemission from Rayleigh by its polarization dependence).

The most confounding case of the model not working as expected was in variations from run-to-run. In Figure B.8, the same sample was measured three times, and there were

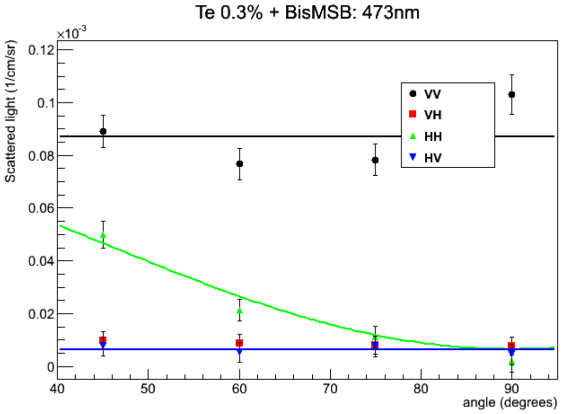

Figure B.7: Data from a scattering measurement.

considerable differences each time. We can see that the measurements can vary by a factor of two or more, run-to-run. We were never able to figure out the reason. The scattering did not trend higher or lower with time, as we might expect if it were becoming contaminated (or if contaminants settled out). The temperature variations in the laboratory should not be able to cause this level of change in scattering nor affect the PMT rate this much. One possibility is that oxygen or water was contaminating the sample, and that the amount of the contaminant depended on the temperature. This could possibly be corrected by bubbling nitrogen though the sample, though we did not have the capability to try this. As the scintillator is an oil, it may not even be possible to do this in a way that does not result in increasing the scattering due to the introduction of bubbles anyway.

B.5 Conclusion

We developed an apparatus to measure the light scattering of samples as a function of wavelength, scattering angle, and polarization. We created a baffle that effectively reduced the background scattering of light. In individual measurements, the scattering fit a model of Rayleigh scattering plus a wavelength-dependent, isotropic absorption-reemission term. However, run-to-run the overall scattering varied considerably for unknown reasons, which

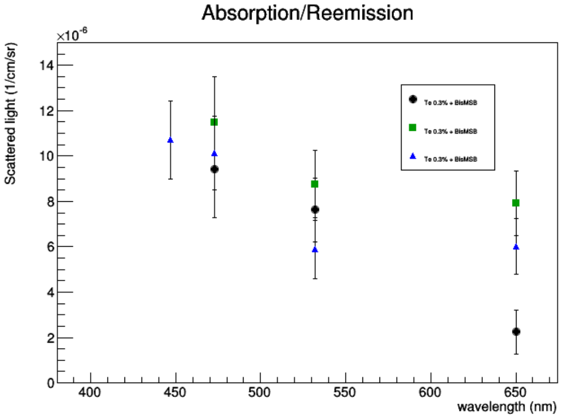

Figure B.8: Fits from three scattering measurements.

ultimately limited the utility of these measurements.