



University of Pennsylvania
ScholarlyCommons

Publicly Accessible Penn Dissertations

1-1-2016

Constraining the Hep Solar Neutrino and Diffuse Supernova Neutrino Background Fluxes With the Sudbury Neutrino Observatory

Andrew Thomas Mastbaum
University of Pennsylvania, amastbaum@gmail.com

Follow this and additional works at: <http://repository.upenn.edu/edissertations>



Part of the [Physics Commons](#)

Recommended Citation

Mastbaum, Andrew Thomas, "Constraining the Hep Solar Neutrino and Diffuse Supernova Neutrino Background Fluxes With the Sudbury Neutrino Observatory" (2016). *Publicly Accessible Penn Dissertations*. 1884.
<http://repository.upenn.edu/edissertations/1884>

Constraining the Hep Solar Neutrino and Diffuse Supernova Neutrino Background Fluxes With the Sudbury Neutrino Observatory

Abstract

The Sudbury Neutrino Observatory has demonstrated that the apparent deficit in solar neutrinos observed on Earth is due to matter-enhanced flavor transitions, and provided precision measurements of the relevant oscillation parameters. The low backgrounds and large, spectral charged-current ν_e -d cross section that enabled these measurements also give SNO unique sensitivity to two yet-unobserved neutrino signals of great interest: the ν_{he} solar neutrino flux and the diffuse supernova neutrino background (DSNB).

This work presents a joint analysis of all three running configurations of the SNO experiment in order to improve constraints on the ν_{he} and DSNB ν_e fluxes. The crucial uncertainties in the energy response and atmospheric neutrino background, as well as the event selection criteria, are reevaluated. Two analysis approaches are taken, a single-bin counting analysis (ν_{he} and DSNB) and multidimensional signal extraction fit (ν_{he}), using a random sample representing 1/3 of the total SNO data. These searches are the most sensitive to date for these important signals, and will improve further when the full dataset is analyzed.

The SNO+ liquid scintillator experiment is a successor to SNO primarily concerned with a search for neutrinoless double-beta decay ($0\nu\beta\beta$) in ^{130}Te . The modifications to the SNO detector in preparation for SNO+ and an analysis of the $0\nu\beta\beta$ sensitivity of this upcoming experiment will also be presented in this work. SNO+ will be the first experiment to load Te into liquid scintillator, and is expected to achieve world-class sensitivity in an initial phase commencing in 2017, with significantly improved sensitivity in an upgraded configuration to follow using much higher Te target mass.

Degree Type

Dissertation

Degree Name

Doctor of Philosophy (PhD)

Graduate Group

Physics & Astronomy

First Advisor

Joshua R. Klein

Subject Categories

Physics

**CONSTRAINING THE HEP SOLAR NEUTRINO AND
DIFFUSE SUPERNOVA NEUTRINO BACKGROUND
FLUXES WITH THE SUDBURY NEUTRINO
OBSERVATORY**

Andrew T. Mastbaum

A DISSERTATION

in

Physics and Astronomy

Presented to the Faculties of the University of Pennsylvania
in

Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

2016

Supervisor of Dissertation

Joshua R. Klein

Professor of Physics

Graduate Group Chairperson

Ravi Sheth

Professor of Physics

Dissertation Committee:

Evelyn Thomson, Associate Professor of Physics

Eugene Beier, Professor of Physics

Masao Sako, Associate Professor of Physics

Adam Lidz, Associate Professor of Physics

**CONSTRAINING THE HEP SOLAR NEUTRINO AND
DIFFUSE SUPERNOVA NEUTRINO BACKGROUND
FLUXES WITH THE SUDBURY NEUTRINO
OBSERVATORY**

COPYRIGHT ©

2016

Andrew T. Mastbaum

This work is licensed under the Creative Commons
Attribution-NonCommercial-ShareAlike 4.0 License.
To view a copy of this license, visit
<http://creativecommons.org/licenses/by-nc-sa/4.0/>

“In cosmological neutrino physics, no neutrinos are harmed. In neutrino astrophysics individual neutrinos must be captured and killed so that we can learn more.”

— John Beacom, *APS April Meeting 2015*

This dissertation is dedicated to all the neutrinos who sacrificed everything
in the name of scientific progress.

Seriously, though, it's for Lauren.

Acknowledgements

None of this work would have been possible without the limitless support of my family, friends, and colleagues. Foremost, I would like to thank my advisor Josh Klein, whose relentless pursuit of the truth in all things has fundamentally shaped the way I approach research. His candor, willingness to entertain (and provide) crazy ideas, and treatment of students as colleagues creates an environment of limitless potential, for which I am extremely grateful, and can only hope to emulate someday. It has been a privilege to work with the legendary Gene Beier, who is always willing to listen and offer wisdom on problems big and small, regarding physics or cats. Bill Heintzelman, with an unwavering commitment to doing things right, has been a source of SNOMAN and *hep* wisdom, and of motivation to always do things right-er. I would like to thank the SNO collaboration as a whole, for building such a phenomenal experiment that is still bearing fruit all these years later.

Penn has been an incredible community of colleagues and friends. Thanks to everyone in the neutrino group past and present who have made DRL a great place to be: Tanner Kaptanoglu, Eric Marzec, Nuno Barros, Ian Coulter, Shannon Glavin, David Rivera, Tom Caldwell, the Rev. Dr. Jeff Secrest. A special thank-you to Richie Bonventre, a “partner in crime” with whom I shared the highest highs and lowest lows of the SNO+ life; Stan Seibert, for whose mentorship and friendship at LANL and Penn I am infinitely grateful; and Gabriel Orebí Gann, a phenomenal scientist whose keen interest in my professional development is always much appreciated. The Penn instrumentation group, especially Rick van Berg and Godwin Mayers, have been wonderful, patient teachers.

I daren’t imagine how different my life might be had I not had the incredible fortune of working at Los Alamos National Laboratory as an undergraduate. I would like to thank my

advisor while at LANL, Andrew Hime, and my wonderful colleagues and mentors. I am so grateful to the Lehigh University Physics Department, especially Michael Stavola, A. Peet Hickman, H. Daniel Ou-Yang, and Gary DeLeo for providing a top-notch education and supporting me in every way possible when my research interests took me out to LANL.

I would also like to thank my family, whose constant encouragement and support has always allowed me to follow my ambitions wherever they might lead. To my parents Tom and Gay, and sister Sara: thank you. To my wonderful cats Chip and Linus: meow meow meow. Finally, to my soon-to-be wife Lauren Cuoco: I will forever be in awe of your patience and grateful for your inexhaustible support as we have navigated this irrationally demanding time in our lives. Now, onward to the great things!

ABSTRACT

CONSTRAINING THE HEP SOLAR NEUTRINO AND DIFFUSE SUPERNOVA NEUTRINO BACKGROUND FLUXES WITH THE SUDBURY NEUTRINO OBSERVATORY

Andrew T. Mastbaum

Joshua R. Klein

The Sudbury Neutrino Observatory has demonstrated that the apparent deficit in solar neutrinos observed on Earth is due to matter-enhanced flavor transitions, and provided precision measurements of the relevant oscillation parameters. The low backgrounds and large, spectral charged-current $\nu_e - d$ cross section that enabled these measurements also give SNO unique sensitivity to two yet-unobserved neutrino signals of great interest: the *hep* solar neutrino flux and the diffuse supernova neutrino background (DSNB).

This work presents a joint analysis of all three running configurations of the SNO experiment in order to improve constraints on the *hep* and DSNB ν_e fluxes. The crucial uncertainties in the energy response and atmospheric neutrino background, as well as the event selection criteria, are reevaluated. Two analysis approaches are taken, a single-bin counting analysis (*hep* and DSNB) and multidimensional signal extraction fit (*hep*), using a random sample representing 1/3 of the total SNO data. These searches are the most sensitive to date for these important signals, and will improve further when the full dataset is analyzed.

The SNO+ liquid scintillator experiment is a successor to SNO primarily concerned with

a search for neutrinoless double-beta decay ($0\nu\beta\beta$) in ^{130}Te . The modifications to the SNO detector in preparation for SNO+ and an analysis of the $0\nu\beta\beta$ sensitivity of this upcoming experiment will also be presented in this work. SNO+ will be the first experiment to load Te into liquid scintillator, and is expected to achieve world-class sensitivity in an initial phase commencing in 2017, with significantly improved sensitivity in an upgraded configuration to follow using much higher Te target mass.

Contents

Title	i
Copyright	ii
Dedication	iii
Acknowledgements	iv
Abstract	vi
Contents	viii
List of Tables	xiii
List of Figures	xvi
1 Introduction	1
2 Physics of Massive Neutrinos	5
2.1 The Standard Model	6
2.2 Neutrino Mass and Mixing	10
2.2.1 Vacuum Oscillations	12
2.2.2 Matter-Enhanced Oscillations	14
2.3 Current Mixing Parameters	17
2.4 Status and Prospects	19

CONTENTS

3 Solar Neutrinos	21
3.1 Standard Solar Models	22
3.2 Solar Neutrino Oscillations	24
3.3 The Homestake Experiment	27
3.4 The Solar Neutrino Problem	29
3.5 Experimental Results	31
3.5.1 Searches for <i>hep</i> Neutrinos	31
4 The Diffuse Supernova Neutrino Background	35
4.1 Core-Collapse Supernovae	36
4.2 The Diffuse Supernova Neutrino Background	38
4.3 Experimental Efforts	41
4.3.1 Current Results	43
4.3.2 Future Directions	44
5 The Sudbury Neutrino Observatory	46
5.1 Physical Principles	47
5.2 Detector	51
5.3 Electronics	53
5.4 Calibration	57
5.4.1 Electronics	57
5.4.2 Sources	60
5.5 Simulation	62
5.5.1 Physics and Detector Modeling	63
5.5.2 Event Rates	64
5.6 Vertex Reconstruction	67
5.6.1 The FTP Vertex Fitter	67
5.6.2 The nFTU Vertex Fitter	68
5.6.3 The RSP Energy Fitter	69
5.6.4 Issues with the FTK Energy Fitter	70
5.7 SNO+ Upgrades	72
5.7.1 Electronics Upgrades	73
6 An Improved <i>hep</i> and DSNB Search with SNO	79
6.1 Signals and Backgrounds	80
6.2 Atmospheric Neutrino Backgrounds	81

CONTENTS

6.2.1	Production	83
6.2.2	Oscillations	85
6.2.3	Atmospheric ν Backgrounds to the <i>hep</i> and DSNB Search	88
6.2.4	Systematic Uncertainties	89
6.2.5	Simulations	90
6.3	Data Selection	90
6.3.1	Run Selection	91
6.3.2	Event Selection	92
6.3.2.1	Low-level Cuts	93
6.3.2.2	High-level Cuts	94
6.4	Blindness	99
6.5	Monte Carlo Simulations	100
6.6	Summary of Systematic Uncertainties	100
6.7	Corrections to Data and Monte Carlo	105
6.8	Counting Analysis Overview	107
6.9	Signal Extraction Overview	108
6.10	Physics Interpretation	110
7	Characterization of Energy-Related Systematic Uncertainties	111
7.1	pT Source Constraints	112
7.2	Michel Electrons	118
7.2.1	Event Selection	118
7.2.2	Energy Fits	121
7.2.3	Parameterization of Systematics	122
7.3	High-Energy Fit Results	123
7.3.1	Ensemble Testing	129
7.4	Constraining Tails: ${}^8\text{Li}$ Source	129
8	Counting Analysis	136
8.1	High-level Cut Optimization	138
8.1.1	Signal Acceptance	142
8.1.2	Summary of Cut Parameters	143
8.2	Low-energy Signal Extraction	146
8.2.1	Ensemble Testing	150
8.3	Energy Window Selection	154
8.4	Sensitivity	160

CONTENTS

8.5	1/3 Dataset Results	163
8.5.1	Projections for the Full Dataset	166
9	Spectral Signal Extraction Fit	172
9.1	Signal Extraction	172
9.2	The Maximum Likelihood Method	173
9.2.1	Floating Systematics	176
9.2.2	Multi-Phase Fitting	176
9.3	Markov Chain Monte Carlo	177
9.3.1	Convergence	180
9.4	Estimating Uncertainty	181
9.4.1	The Bayesian Interpretation	183
9.4.2	The Profile Likelihood Construction	184
9.4.3	A Two-dimensional Example	186
9.5	The <code>sxmc</code> Code	189
9.5.1	Implementation	190
9.5.2	Proposal Distributions and Burn-in	191
9.6	Signals and Backgrounds	192
9.6.1	Cuts	193
9.6.2	Observables	194
9.6.3	Systematic Uncertainties	196
9.7	Ensemble Testing	197
9.8	Results	198
9.8.1	Sensitivity	198
9.8.2	1/3 Data Results	200
9.8.3	Goodness of Fit	205
10	Conclusions	210
A	The SNO+ Experiment	214
A.1	Neutrinoless Double Beta Decay	215
A.1.1	Current Results	220
A.2	Backgrounds	222
A.3	Counting Analysis	223
A.3.1	Sensitivity	224
A.3.2	Reconstruction	225

CONTENTS

A.3.3	Cuts	226
A.3.4	Optimizing the Signal Region of Interest	227
A.3.5	Results	228
A.3.6	Systematic Uncertainties	232
A.4	SNO+ Phase II	233
A.5	Conclusions	235
B	Atmospheric Neutrino Event Simulation with GENIE	238
B.1	Physics Models	239
B.2	Detector Model	239
C	A Time-Correlated Approach to Supernova Background Detection	241
D	Low-Level Cuts	244
E	pT Source Run List	246
F	${}^8\text{Li}$ Source Run List	247
G	Frequentism	249
References		251

List of Tables

2.1	Matter content in the Standard Model. For each particle there exists a corresponding antiparticle having identical mass but opposite charges.	7
2.2	Three-flavor oscillation parameters from our fit to global data after the NOW 2014 conference. The results are presented for the “Free Fluxes + RSBL” in which reactor fluxes have been left free in the fit and short baseline reactor data (RSBL) with $L \lesssim 100$ m are included. The numbers in the 1st (2nd) column are obtained assuming NO (IO), <i>i.e.</i> , relative to the respective local minimum, whereas in the 3rd column we minimize also with respect to the ordering. Note that $\Delta m_{3\ell}^2 \equiv \Delta m_{31}^2 > 0$ for NO and $\Delta m_{3\ell}^2 \equiv \Delta m_{32}^2 < 0$ for IO. Table and caption reproduced from [29]. “bfp” refers to the best-fit parameter value.	18
3.1	Solar neutrino fluxes in $\text{cm}^{-1} \text{ s}^{-1}$ for a variety of SSMS. Table reproduced from Reference [21].	24
3.2	Summary of solar neutrino flux measurements.	32
4.1	Cosmological parameters as measured by Planck in [64]. H_0 is given in units of $\text{km s}^{-1} \text{ Mpc}^{-1}$. See reference for further details on the different constraints.	39
5.1	Neutrino interactions in SNO.	49
6.1	Total fluxes of atmospheric neutrinos predicted in the Bartol04 model, integrated over energy and zenith angle.	85
6.2	Total oscillated fluxes of atmospheric neutrinos predicted by the Bartol04 model and the oscillation model described in the text, integrated over energy and zenith angle.	87

LIST OF TABLES

6.3	Expected number of atmospheric neutrino events per year in the heavy water during Phase I.	87
6.4	Systematic uncertainties. “Correlated” means that the parameter is assumed to take on the same value across all phases.	101
7.1	Results of pT source fits. Units are MeV, except shifts which are fractional.	118
7.2	Results of lifetime fits to the selected decay electron events.	119
7.3	Results of the Michel electron fits.	125
7.4	Constraints used in the Michel electron energy systematics fit, and the fit results evaluated at the constraint energies.	126
7.5	Correlation coefficients for Phase I Michel fit.	126
7.6	Results of direct fits to the source data shown in Figure 7.5.	126
8.1	High-level cuts for each region of interest.	146
8.2	Raw numbers events in the Phase I high-energy atmospheric MC with $14 < T_{\text{eff}} < 20$ MeV and $r < 550$ cm passing each cut. Note that as these numbers are for the raw MC production, and initial state neutrino flavor ratios are not correct.	147
8.3	Results of the signal extraction fits with 1/3 data, in terms of the fractional scaling from the model prediction (using BS05(OP) fluxes).	150
8.4	Results of the counting experiment with the 1/3 data set, with all systematic parameters fixed to their mean values.	162
8.5	Characteristics of observed events in the 1/3 data set. r_{NCD} is the distance to the closest NCD.	162
8.6	Summary of systematics sampling results for the counting analysis signal and background expectations for the 1/3 data set.	170
8.7	Summary of the 1σ uncertainties on the signal (ΔS) and background (ΔB) expectations in the <i>hep</i> and DSNB ROI in each phase due to the variation of a single systematic parameter while others are fixed to their mean value. All units are percentages ($\Delta A = 100.0 \times \sigma_A/\mu_A$).	171
9.1	Performance of sxmc PDF histogram building on a variety of platforms. . .	191
9.2	Best-fit values with 90% CL profile likelihood errors and 1σ constraints for the parameters floated in the joint three-phase signal extraction fit to the 1/3 dataset, with scanned systematics at their nominal values. Flux scales here are relative to the BP2000 model prediction.	203

LIST OF TABLES

9.3	Upper limits in units of the BS05(OP) SSM prediction, obtained with scanned systematics parameters individually varied in the joint three-phase signal extraction fit to the 1/3 dataset.	204
A.1	Expected background counts in the ROI in the first year of SNO+ data taking.	230
A.2	Expected sensitivity of the SNO+ experiment at the 90% confidence level. See note in the text about the assumptions regarding the matrix elements used for the $m_{\beta\beta}$ estimate.	231
A.3	Effective Majorana mass limits for a nominal five-year SNO+ run, for a variety of theoretical approaches to matrix element calculation.	231
A.4	Systematic parameters and constraints.	234
A.5	Projections for SNO+ Phases I and II compared, each for a nominal five-year live time.	235
C.1	Number of time-integrated neutrino events in each SNO phase for dominant channels, above a threshold of 16 hit PMTs, according to the Beacom and Vogel model [165]. Data from Reference [164], Table 7.1.	242

List of Figures

2.1	Charged-current $\nu_e - e^-$ scattering, which results in the resonant conversion of ν_e to other flavors in the MSW effect.	15
3.1	The reactions of the solar pp fusion chain, with the names of the neutrino-producing reactions are given in parentheses. Figure based on Fig. 10.1 of Reference [21].	23
3.2	Solar neutrino spectrum. This figure shows the energy spectrum of neutrinos predicted by the standard solar model (Bahcall and Pinsonneault 2004). The neutrino fluxes from continuum sources (like pp and 8B) are given in the units of number per cm^2 per second per MeV at one AU. The line fluxes (pep and ^7Be) are given in number per cm^2 per second. The spectra from the pp chain are drawn with solid lines; the neutrino energy spectra from reactions with carbon, nitrogen, and oxygen (CNO) isotopes are drawn with dotted lines. Figure from Reference [33].	25
3.3	Spatial distributions of the fluxes in the GS98SFII standard solar model. . .	26
3.4	The survival probability P_{ee} for 8B and hep solar neutrinos, assuming the BS05(OP) SSM, and PDG-recommended three-neutrino mixing parameters (with their uncertainties shown as bands), according to the Sun-Earth Large Mixing Angle Adiabatic (SELMAA) calculation (see Reference [38]).	28
3.5	The initial Super-Kamiokande data showing a potential hep excess. Figures from Reference [59].	33
3.6	The published 2001 Super-Kamiokande data showing a hep flux consistent with the SSM. Figures from Reference [60].	33

LIST OF FIGURES

3.7 The spectrum published by SNO for events near the ${}^8\text{B}$ endpoint, indicating the <i>hep</i> search region. Shaded bands show the systematic uncertainties. Figure from Reference [15].	34
4.1 Comparison of data with the Strigari concordance model for the star formation rate. The hatched band is allowed by observations including the SuperK DSNB $\bar{\nu}_e$ limit. The black, red, and blue lines are measurements by Dahlen et al. [66], GALEX [67], and 2MASS+2dF [68], respectively, and indicating the effect of a dust correction for the latter. Figure from [65].	40
4.2 Expected flux of DSNB neutrinos on Earth, assuming a thermal spectrum. The unoscillated solar ${}^8\text{B}$ and <i>hep</i> fluxes are shown for scale. Compare to Figure 1 in [63].	42
5.1 Normalized differential cross sections for the $(\nu_e - d)$ CC and $(\nu_e - e^-)$ ES reactions, showing the spectral fidelity of CC and directionality of the ES. Figures reproduced from Reference [27].	48
5.2 Schematic diagrams illustrating the neutrino interactions in SNO.	49
5.3 The SNO Experiment. Figures from Reference [77].	52
5.4 A simplified overview of the SNO data acquisition electronics. Figure from Reference [77].	55
5.5 An overview of the SNO trigger system. Figure from Reference [78].	58
5.6 The SNO source manipulator system. Figure from Reference [77].	60
5.7 Total cross sections for $\nu_e - d$ CC, $\nu - d$ NC, and $\nu - e^-$ ES.	66
5.8 Examples of FTK mis-reconstruction of high-energy events, shown for Phase I isotropic electron Monte Carlo after corrections.	71
5.9 A comparison of SNO fitter performance for reconstructing Phase I stopped μ Monte Carlo.	72
5.10 The SNO+ detector. The active volume is contained within a 12 m diameter acrylic sphere (grey), suspended with ropes in a volume of ultra-pure water. The outer 18 m diameter structure supports 9500 inward-looking PMTs.	73
5.11 A block-level overview of the major features of the MTC/A+, the upgraded analog trigger board for SNO+.	76
6.1 Energy spectra around the <i>hep</i> and DSNB energy regions of interest for Phase I, after corrections and cuts described in the text.	82

LIST OF FIGURES

6.2 Fluxes calculated in the Bartol04 calculation for each flavor, averaged over zenith angle. Min and max refer to the solar minimum and maximum. The summed spectra for $\nu_e + \bar{\nu}_e$ and $\nu_\mu + \bar{\nu}_\mu$ at the solar minimum are shown for comparison to Figure 2 in [96].	84
6.3 Energy spectra for low-energy ($E_\nu < 100$ MeV) atmospheric neutrino CC interactions in Phase I.	89
6.4 Signal sacrifice due to instrumental background cuts, as measured in the Phase I [103]	104
6.5 The shape uncertainty in the Winter solar ${}^8\text{B}$ ν_e spectrum.	105
7.1 Distributions of pT source data at the origin.	114
7.2 Energy distributions of pT source data at the origin, demonstrating the background reduction due to the cuts and the Gaussian fit to extract the mean and resolution.	115
7.3 pT source data (black) and Monte Carlo (red) comparisons as a function of radius, and the corresponding shifts.	117
7.4 Reconstructed energy spectra and time differences between predecessor and retrigger events for the Michel electron event selection.	120
7.5 Results of independent source fits for energy systematics.	124
7.6 Best fit to Michel electrons for energy systematics.	127
7.7 Energy dependence of the scale parameters extracted in the Michel electron fit.	128
7.8 Ensemble testing distributions for the Phase I energy systematics fit using Michel electrons.	130
7.9 Summary of pull distributions for energy-related systematic parameters.	131
7.10 Cut parameter distributions for Phase I central ${}^8\text{Li}$ source run 14348.	132
7.11 Comparison of data (black points) to Monte Carlo (gray bands) for ${}^8\text{Li}$ source runs. The MC bands indicate systematic uncertainties due to energy scale and resolution as well as MC statistics. Source deployment positions are given in Appendix F.	134
8.1 Examples of atmospheric neutrino Monte Carlo events removed by the high-level cuts. The color indicates time a PMT was hit: green is earlier, red is later.	139

LIST OF FIGURES

8.2	Data and Monte Carlo distributions for the 1D angular KS test (P_ϕ) high-level cut parameter in Phase I. The middle and bottom panels show the leakage fraction of atmospheric neutrino background and the sacrifice of the signal as a function of cut value.	140
8.3	High-level cut optimization scans for the <i>hep</i> ROI in Phase I.	141
8.4	Discrepancy between the data and MC for the sacrifice of the level-cuts due to ITR. Phase II data is shown.	143
8.5	Data/Monte Carlo comparison for the total signal sacrifice of the high-level cuts in each Phase.	144
8.6	Sacrifice in Monte Carlo electrons due to high-level cuts. Discontinuities are due to the three different ROIs which in general have different cuts applied.	145
8.7	Effective electron-equivalent kinetic energy spectra for atmospheric neutrino backgrounds for Phase I after application of the cuts summarized in Tables 8.2 and 8.1.	148
8.8	Results of the Phase I signal extraction fit with 1/3 data.	151
8.9	Results of the Phase II signal extraction fit with 1/3 data.	152
8.10	Results of the Phase III signal extraction fit with 1/3 data.	153
8.11	Ensemble tests for the Phase I low-energy signal extraction.	155
8.12	Summary of pull distributions for low-energy signal extraction parameters.	156
8.13	The ROI optimization for the Phase I.	157
8.14	The ROI optimization for the Phase II.	158
8.15	The ROI optimization for the Phase III.	159
8.16	Energy spectra around the <i>hep</i> and DSNB regions of interest, after corrections and all cuts except on energy, with all parameters fixed to their mean values.	161
8.17	Reconstructed positions and directions of events observed in the <i>hep</i> ROI in the 1/3 dataset ($x - y$ projection). NCDs, shown as gray circles, are enlarged by a factor of 5. The arrows (which are arbitrarily but uniformly scaled) represent the reconstructed event direction, projected onto the $x - y$ plane. Further details about these events are listed in Table 8.5.	165
8.18	Monte Carlo systematic parameter variation in Phase I.	167
8.19	Monte Carlo systematic parameter variation in Phase II.	168
8.20	Monte Carlo systematic parameter variation in Phase III.	169
9.1	A three-state system.	178

LIST OF FIGURES

9.2	Examples of parameter time series and projected distributions for a poorly-mixed and well-mixed chain.	182
9.3	A fit to a signal-free fake experiment in a two-signal example.	187
9.4	A comparison of sensitivities obtained by integration and the Feldman-Cousins approach.	188
9.5	One-dimensional projections of the PDFs for Phase I.	195
9.6	Coverage of 90% CL intervals in fake data ensemble tests using a variety of true <i>hep</i> fluxes. Error bars are statistical, due to the finite size of the ensemble. Thick lines are for fits with fixed systematics and thin lines for fits with floating systematics.	199
9.7	Marginalized posterior distribution for the <i>hep</i> flux, used to determine credible intervals. The 90% CL interval is highlighted in gray.	201
9.8	Projections of the best fit to the 1/3 dataset. The <i>hep</i> CC and ES signals are shown at the 90% CL upper limit. Phases I, II, and III are shown in the left, center, and right columns, respectively.	202
9.9	Fits to shifted systematic parameters.	206
9.10	Monte Carlo-derived probability distribution for the Pearson χ^2 statistic.	208
9.11	Parameter values as a function of time in the MCMC likelihood evaluation.	209
A.1	Double-beta decay spectra, including a few-percent energy resolution.	217
A.2	Light Majorana neutrino exchange, one possible mechanism for $0\nu\beta\beta$	217
A.3	A Vissani-Strumia plot showing the parameter space for $0\nu\beta\beta$ in the canonical three-neutrino mixing model.	219
A.4	Grid scan of the joint fiducial volume/energy range parameter space to optimize the signal ROI in the SNO+ counting experiment. Sensitivity is approximate only.	228
A.5	Expected backgrounds with $r < 3.5$ m, shown with an example signal for a Majorana neutrino with $m_{\beta\beta} = 200$ meV.	229
A.6	A Vissani-Strumia plot showing the parameter space for $0\nu\beta\beta$ in the canonical three-neutrino mixing model, with the projected sensitivity for SNO+ as well as selected experimental limits (see Section A.1.1 for details).	237
B.1	Visualization of the simplified SNO geometry used in the GENIE primary vertex generation. The detector model includes the acrylic vessel with neck (blue), light water (gray), and heavy water. The facets are an artifact of the visualization.	240

Chapter 1

Introduction

Decades of neutrino experiments using solar [1–7] and terrestrial [8–11] neutrino sources have firmly established that neutrinos undergo flavor-changing oscillations and have small but nonzero mass, and demonstrated the validity of solar energy production models. Of the several electron neutrino-producing reactions in the Sun, the highest-energy *hep* neutrinos remain unobserved. A measurement of this flux would provide a verification of solar energy production models complementary to recent low-energy results. The *hep* neutrinos also have unique sensitivity to certain classes of non-standard solar models, as they are produced at a relatively high energy and high radius within the Sun [12]. The heavy water (${}^2\text{H}_2\text{O}$, hereafter “D₂O”)-based Sudbury Neutrino Observatory experiment (SNO) had unique sensitivity to low-energy ν_e including *hep* solar neutrinos through the large charged-current (CC) cross section on deuterium, which is about 100 times larger than the $\nu_e - e^-$ elastic scattering (ES) cross section at 20 MeV. The CC interaction also provides a better measurement of the incoming neutrino energy, as the final state electron kinetic energy is strongly correlated to the initial state neutrino energy. These benefits allow SNO to reach a sensitivity beyond

even much larger light water Cherenkov detectors.

According to astrophysical models of core-collapse supernova, the high frequency of these events within the observable universe produces an MeV-scale neutrino and antineutrino “glow” that should be observable in terrestrial neutrino observatories. This too remains unobserved so far. Measurement of this diffuse supernova neutrino background (DSNB) would provide valuable input to our understanding of supernovae, as it probes both the average number of neutrinos released in these events and the average temperature at the surface of last scattering for the neutrinos. In contrast to single nearby core-collapse supernova events like SN1987A [13], the DSNB is always “on,” and the information about average behavior provides context for interpreting isolated nearby supernova events. The Super-Kamiokande experiment, using a larger, light water Cherenkov detector, has set limits on the DSNB $\bar{\nu}_e$ flux via inverse beta decay searches which are close to the model predictions [14]. The SNO detector, on the other hand, had world-class sensitivity to the DSNB ν_e flux, again thanks to the large, spectral CC $\nu_e - d$ cross section and to very low backgrounds.

Following the first phase of the SNO experiment with a D₂O target, the SNO collaboration published limits on the *hep* solar neutrino flux and the ν_e component of the DSNB flux in a combined analysis [15], and these remain the most stringent direct limits today. Subsequent to the collection of that data, the SNO detector was run in two more configurations, which differed primarily in the mechanism for neutron detection, for a total of three “phases.” For the second phase, a small concentration of NaCl was added to the D₂O, and in the third phase, the NaCl was removed and an array of proportional counters was deployed in the D₂O. Studies by C. Howard in 2010 investigated a combined analysis of all three phases to search for the *hep* flux only, using a multidimensional fit [16]. That work suggests that a low-significance measurement of the *hep* flux is possible using only 1/3 of the full SNO dataset, but did not fully account for systematic uncertainties (most notably the crucial energy-related uncertainties above the endpoint of the ⁸B solar neutrino spectrum)

or possible biases in the result. This “hint” of a detection provides strong motivation for a thorough independent analysis, which I undertake here.

This dissertation expands on the previous work, and uses data from all three phases of the SNO experiment to constrain both the *hep* and DSNB ν_e fluxes. I perform the first counting analysis of the combined (three phase) data, in order to produce both limits which may be directly compared to the published results and model-independent counting rates. This work represents the sole effort toward improving SNO’s (world-leading) direct DSNB ν_e flux limit. I also conduct a multidimensional spectral fit for the *hep* flux, which has been newly developed for this work and is substantially more sophisticated than previous efforts, particularly in the statistical treatment. Both the counting and signal extraction approaches benefit from the higher statistics and better measurement of correlated backgrounds afforded by the expanded data set, and also take advantage of more recent theoretical progress and analysis improvements. I have also re-evaluated the energy-related systematic uncertainties, using calibration data to define and validate an energy response model covering the full analysis energy range, and improved the modeling of atmospheric neutrino interactions, revisiting the dominant uncertainties for these searches to ensure a robust result.

I begin this dissertation with brief introductions to the relevant physics: the Standard Model (Chapter 2); solar neutrinos (Chapter 3); and the diffuse supernova neutrino background (Chapter 4). Next, the SNO detector is introduced, along with SNO+-related upgrades (Chapter 5), and the reconstruction of events is briefly reviewed. Chapter 6 presents an overview of the *hep* and DSNB searches, including the signals and backgrounds, data selection process, systematic uncertainties, and details on atmospheric neutrino backgrounds. The measurement of the systematic uncertainties associated with energy reconstruction, which are of great importance for the *hep* measurement in particular, is detailed in Chapter 7. A counting (“box”) analysis approach to the *hep* and DSNB measurement is presented in Chapter 8, which details the selection of data, the optimization of cuts, and the statistical treatment. Next, a multidimensional spectral signal fit to extract the *hep* is presented in

Chapter 9. Finally Chapter 10 provides a summary and conclusions.

Since the conclusion of the SNO experiment, a new experiment known as SNO+ has been under development that will reuse much of the SNO detector infrastructure. SNO+ will use a liquid scintillator, rather than heavy water, target to explore a broad range of important questions in neutrino physics, with the primary objective being a sensitive search for neutrinoless double-beta decay ($0\nu\beta\beta$) in ^{130}Te . The upgrades to the detector and the $0\nu\beta\beta$ sensitivity analysis (which has strong parallels to the *hep* search in SNO) are presented in Appendix A.

Chapter 2

Physics of Massive Neutrinos

A light, neutral, weakly-interacting “neutron” was first proposed by Pauli in 1930 [17] in order to explain the apparent energy nonconservation observed in β decays. The ghostly particle smuggling away the energy was soon renamed the “neutrino” (the little neutral one) in 1933 by Fermi to avoid confusion with the modern neutron discovered by Chadwick in 1932. In the intervening years, enormous effort has gone into quantifying “little” and “neutral.”

The neutrino was first observed experimentally in 1956 by Cowan and Reines through the inverse beta decay of antineutrinos ($\bar{\nu}$) from the Savannah River nuclear reactor [18]. Subsequent experiments established that (left-handed) neutrinos and (right-handed) antineutrinos are distinct, and that three independent neutrino flavors ensure lepton flavor conservation in weak interactions. With the discovery of neutrino flavor oscillations in atmospheric neutrinos in 1998 by the Super-Kamiokande Collaboration [19] and in solar neutrinos in 2001 by the SNO Collaboration [6], however, came the implication of nonzero neutrino mass and lepton flavor nonconservation, opening up rich new phenomenology in

the lepton sector. In some sense, less is “known” about the neutrino today than in 1930!

2.1 The Standard Model

The Standard Model developed in the 1970s brought order to a proliferation of newly discovered particles. It described the particles comprising the known matter of the universe and provided a unified description of their strong, weak, and electromagnetic interactions. Notably, the Standard Model did not include gravitation, and from a modern perspective, was lacking in its exclusion of more recently discovered neutrino masses, dark matter, and dark energy. Although neutrino masses have since been accommodated, the remaining limitations strongly suggest that the Standard Model is an effective theory that is part of some more comprehensive model. This has motivated an enormous body of theoretical work and experimental searches for “beyond the Standard Model” particles and interactions.

The Standard Model is a quantum field theory that asserts gauge invariance of the composite symmetry group

$$\mathrm{SU}(3)_C \times \mathrm{SU}(2)_L \times \mathrm{U}(1)_Y,$$

with C referring to color, L to left-handed chirality, and Y to weak hypercharge. Gauge invariance demands the introduction of massless bosons corresponding to the generators of each group, which define the interactions between matter particles. These are denoted $g_{1..8}$, $A_{1..3}^\mu$, and B^μ . The gluons, g , mediate the strong nuclear force and are described by quantum chromodynamics (QCD), while the A and B together mediate weak nuclear and electromagnetic interactions. The strengths of these interactions are given by coupling constants, which are free parameters of the theory and must be determined experimentally.

The matter content consists of fermionic fields, which are expressed in representations of the groups. These are divided into *quarks* which partipate in all three types of interactions, and *leptons* which feel only the weak and electromagnetic forces. In nature, there exist

		Generation		
		1	2	3
quarks		u, d	c, s	t, b
leptons		ν_e, e	ν_μ, μ	ν_τ, τ

Table 2.1: Matter content in the Standard Model. For each particle there exists a corresponding antiparticle having identical mass but opposite charges.

three generations of matter, which share quantum numbers but differ in mass. The matter content of the Standard Model is summarized in Table 2.1. In the chiral representation conventionally used for analysis of highly relativistic neutrinos, the quarks and leptons are represented as:

$$L_\alpha = \begin{pmatrix} \nu_{Li} \\ \ell_{Li} \end{pmatrix}, Q_\alpha = \begin{pmatrix} U_{Li} \\ D_{Li} \end{pmatrix}, \ell_{Ri}, \nu_{Ri}, U_{Ri}, D_{Ri}, \quad (2.1)$$

with L and Q representing the lepton and quark doublets; ℓ the massive leptons; U and D the up- and down-type quarks, respectively; subscripts L and R for left- and right-handed chirality; and i the generation.

At the time the Standard Model was developed, experimental evidence suggested maximal parity violation in neutrino interactions, with only left-handed chiral neutrinos and right-handed antineutrinos participating. Therefore the right-handed neutrino, a non-interacting $SU(2)_L$ singlet, was excluded from the theory.

According to the electroweak theory developed by Weinberg, Glashow, and Salam (see, e.g., Reference [20]), the electromagnetic and weak interactions are combined by mixing of the gauge bosons A and B . Specifically, the trilinear couplings between electrically-charged

2.1 The Standard Model

particles and photons take the form¹

$$\mathcal{L}_{\text{int}}^{(\gamma)} = -\ell j_\gamma^\mu A_\mu \quad (2.2)$$

$$j_\gamma^\mu = -\bar{\ell} \gamma^\mu \ell \quad (2.3)$$

and new, mixed fields are formed as

$$W^\mu = \frac{A_1^\mu - iA_2^\mu}{\sqrt{2}} \quad (2.4)$$

$$A^\mu = \sin \theta A_3^\mu + \cos \theta B^\mu \quad (2.5)$$

$$Z^\mu = \cos \theta A_3^\mu - \sin \theta B^\mu \quad (2.6)$$

where θ is the weak mixing angle. These W and Z are the vector bosons that participate in charged-current (CC) and neutral-current (NC) weak interactions, respectively. In the interacting Lagrangian, these take the forms

$$\mathcal{L}_{\text{int}}^{(\text{CC})} = -\frac{g_W}{2\sqrt{2}} j_{W,L}^\mu W_\mu + h.c. \quad (2.7)$$

$$j_{W,L}^\mu = \bar{\nu}_\ell \gamma^\mu (1 - \gamma^5) \ell = \bar{\nu}_{\ell L} \gamma^\mu \ell_L \quad (2.8)$$

$$\mathcal{L}_{\text{int}}^{(\text{NC})} = -\frac{g_W}{2 \cos \theta} j_{Z,L}^\mu Z_\mu + \mathcal{L}_{\text{int}}^{(\gamma)} \quad (2.9)$$

$$j_{Z,L}^\mu = 2g_L^\nu \bar{\nu}_{\ell L} \gamma^\mu \nu_{\ell L} + 2g_L^l \bar{\ell}_L \gamma^\mu \ell_L + 2g_R^l \bar{\ell}_R \gamma^\mu \ell_R. \quad (2.10)$$

¹This section largely follows the notation of [21].

2.1 The Standard Model

Unlike its CC counterpart, the NC interaction is diagonal in the flavor basis, and so neutral current interactions cannot change lepton flavor.

The W and Z are known experimentally to have nonzero mass, which the Standard Model accounts for via the so-called Higgs mechanism. In this scheme, a new $SU(2)_L$ doublet is introduced which is composed of a charged and neutral complex scalar field:

$$\Phi(x) = \begin{pmatrix} \phi^+(x) \\ \phi^0(x) \end{pmatrix}. \quad (2.11)$$

The Higgs term in the Lagrangian takes the form

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi) \quad (2.12)$$

with the potential term

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2 = \lambda \left(\Phi^\dagger \Phi - \frac{v}{2} \right)^2 \quad (2.13)$$

where $v = \sqrt{\mu^2/\lambda}$, $\mu < 0$, and $\lambda > 0$. This potential has a minimum at $\Phi^\dagger \Phi = v^2/2$, which corresponds to the minimum-energy (vacuum) state for the field. This implies that the charged field must have zero expectation value at the vacuum while the neutral part must be finite:

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \quad (2.14)$$

This “breaks” the $SU(2) \times U(1)$ symmetry but preserves $U(1)_Q$, which conserves electric

2.2 Neutrino Mass and Mixing

charge. Working in the unitary gauge where

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (2.15)$$

and expanding the covariant derivatives in Equation 2.12, one finds, in addition to kinetic and mass terms for the Higgs boson, interaction terms with the W and Z and mass-like terms involving WH and ZH interactions. It is through the coupling to the Higgs field that the vector bosons of the electroweak interaction gain their observed masses.

The Higgs field also couples to fermions through Yukawa couplings, generating the fermion masses:

$$\mathcal{L}_{HL} = - \sum_{\alpha} \frac{y_{\alpha}^l v}{\sqrt{2}} \bar{\ell}_{\alpha} \ell_{\alpha} - \sum_{\alpha} \frac{y_{\alpha}^l}{\sqrt{2}} \bar{\ell}_{\alpha} \ell_{\alpha} H \quad (2.16)$$

where $\ell_{\alpha} = \ell_{\alpha L} + \ell_{\alpha R}$ and $\alpha = e, \mu, \tau$, and y are Yukawa coupling constants. The fermion masses, in terms of the unknown couplings and the Higgs vacuum expectation value (VEV) v , are thus given by

$$m_{\alpha} = \frac{y_{\alpha}^l v}{\sqrt{2}}. \quad (2.17)$$

A similar process generates masses for the quarks.

The Higgs boson, with a mass of $125 \text{ GeV}/c^2$, was discovered simultaneously in 2012 by the ATLAS [22] and CMS [23] collaborations at the Large Hadron Collider, a monumental feat which has finally validated this linchpin of the Standard Model.

2.2 Neutrino Mass and Mixing

The success of the electroweak theory as originally envisioned was relatively short-lived, as the observation of neutrino oscillations in 1998 [19] implied that neutrinos have mass (and

2.2 Neutrino Mass and Mixing

ν_R is required in the theory) and that lepton flavor is not conserved.

Including the right-handed singlets, new terms are allowed in the Lagrangian:

$$\mathcal{L}_D = -m_{ij}^\nu \overline{\nu_{Ri}} \nu_{Lj} \quad (2.18)$$

$$\mathcal{L}_M = -M_{ij}^\nu \overline{\nu_{Ri}} \nu_{Rj}^c. \quad (2.19)$$

The former is the typical Dirac mass term which exists for all fermions, and the latter a Majorana mass term, only allowed for the neutral neutrino. If the neutrino is a Majorana fermion, it is its own antiparticle in the sense that the distinction between “neutrinos” and “antineutrinos” is fundamentally a statement about chirality states of the same particle. This is a question of profound significance, but has no observable impact on the phenomenology of neutrino oscillations and may be neglected in the following discussions of solar neutrinos. Determination of the (Majorana or Dirac) nature of the neutrino, however, is an area of intense research interest, in particular through searches for neutrinoless double beta decay ($0\nu\beta\beta$); Appendix A describes such a search with the SNO+ experiment, which is built on the infrastructure and experience of the SNO experiment.

So far, the mass matrix in \mathcal{L}_D has been assumed to be diagonal, i.e. flavor-conserving. In general, however,

$$\mathcal{L}_D = - \begin{pmatrix} \overline{\nu_{eR}} & \overline{\nu_{\mu R}} & \overline{\nu_{\tau R}} \end{pmatrix} \begin{pmatrix} m_{ee} & m_{e\mu} & m_{e\tau} \\ m_{\mu e} & m_{\mu\mu} & m_{\mu\tau} \\ m_{\tau e} & m_{\tau\mu} & m_{\tau\tau} \end{pmatrix} \begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} + h.c. \quad (2.20)$$

$$\mathcal{L}_D = - \begin{pmatrix} \overline{\nu_{eR}} & \overline{\nu_{\mu R}} & \overline{\nu_{\tau R}} \end{pmatrix} U \begin{pmatrix} m_1 & 0 & 0 \\ 0 & m_2 & 0 \\ 0 & 0 & m_3 \end{pmatrix} U^{-1} \begin{pmatrix} \nu_{eL} \\ \nu_{\mu L} \\ \nu_{\tau L} \end{pmatrix} + h.c. \quad (2.21)$$

2.2 Neutrino Mass and Mixing

where U is the unitary Pontecorvo-Maki-Nakagawa-Sakata, or PMNS, matrix which diagonalizes m_{ij} . It is conventional to parameterize this matrix in terms of three mixing angles θ_{12} , θ_{23} , and θ_{13} , and a complex phase δ , separated by factorization²:

$$U_{\text{PMNS}} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \times \begin{pmatrix} c_{13} & 0 & s_{13}e^{i\delta} \\ 0 & 0 & 0 \\ -s_{13}e^{-i\delta} & 0 & c_{13} \end{pmatrix} \times \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (2.22)$$

with $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. This representation isolates the parameters typically probed by different types of experiments — the “atmospheric” mixing angle θ_{23} , the “solar” θ_{12} , and the “reactor/accelerator” θ_{13} — so differentiated for reasons that will be described shortly.

The PMNS matrix provides a transformation between the flavor eigenstates of the weak interaction and the mass eigenstates of the free-particle Hamiltonian. Therefore, we may expand a flavor state $\alpha = e, \mu, \tau$ in the basis of mass eigenstates $i = 1, 2, 3$:

$$|\nu_\alpha\rangle = \sum_i U_{\alpha i}^* |\nu_i\rangle. \quad (2.23)$$

2.2.1 Vacuum Oscillations

Applying Schrödinger’s equation, one finds that the probability of transition between two flavor eigenstates α and β at time t is given by

$$P_{\alpha\beta} = P(\alpha \rightarrow \beta) = |\langle \nu_\alpha | \nu_\beta \rangle|^2 = \sum_{j,k} U_{\alpha j} U_{\beta j}^* U_{\alpha k}^* U_{\beta k} e^{-i(E_j - E_k)t}. \quad (2.24)$$

²In the Majorana case, additional degrees of freedom require the addition of two more phases. This is typically parameterized as $U'_{\text{PMNS}} = U_{\text{PMNS}} \times \text{diag}(1, e^{i\phi_1}, e^{i\phi_2})$.

2.2 Neutrino Mass and Mixing

This may be re-expressed in terms of the mass difference squared between the mass eigenstates, recognizing that for the highly relativistic neutrinos, time $t \sim$ position L and the mass/momentum ratio $\ll 1$:

$$P(\alpha \rightarrow \beta) = \delta_{\alpha\beta} - 4 \sum_{i>j} \Re(U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}) \sin^2(1.27 \Delta m_{ij}^2 (L/E)) \quad (2.25)$$

$$+ 2 \sum_{i<j} \Im(U_{\alpha i} U_{\beta i}^* U_{\alpha j}^* U_{\beta j}) \sin^2(2.54 \Delta m_{ij}^2 (L/E)) \quad (2.26)$$

with L the distance in km, E the energy in GeV, and $\Delta m_{ij}^2 \equiv m_i^2 - m_j^2$.

Considering only two neutrino flavors, the mixing matrix has only one free parameter, a mixing angle θ_V . This simplifying assumption is a good approximation for solar neutrino experiments, where states $|\nu_e\rangle$ and $|\nu_a\rangle$ (a linear combination of $|\nu_\mu\rangle$ and $|\nu_\tau\rangle$) are considered — the latter cannot be distinguished because solar neutrinos lack the energy to produce the associated charged leptons. In the quasi-two neutrino case,

$$U = \begin{pmatrix} \cos \theta_V & \sin \theta_V \\ -\sin \theta_V & \cos \theta_V \end{pmatrix}. \quad (2.27)$$

Hence, for a ν_e ,

$$|\nu_e(x, t)\rangle = \cos \theta_V e^{-ip_1 x} |\nu_1\rangle + \sin \theta_V e^{-ip_2 x} |\nu_2\rangle. \quad (2.28)$$

Inverting the expressions for $|\nu_e\rangle$ and $|\nu_a\rangle$ to write the free-particle Hamiltonian in terms of flavor eigenstates:

$$\mathcal{H} = \frac{m_1^2 + m_2^2}{4E} + \frac{\Delta m_{21}^2}{4E} \begin{pmatrix} -\cos 2\theta & \sin 2\theta \\ \sin 2\theta & \cos 2\theta \end{pmatrix}. \quad (2.29)$$

2.2 Neutrino Mass and Mixing

At a time t and distance L , $p_i x = E_i t - p_i L$. Assuming that $p_1 = p_2$ (hence $E_1 \neq E_2$)³,

$$\langle \nu_e | \nu_e \rangle = \cos^2 \theta_V e^{-eE_1 t} + \sin^2 \theta_V e^{-iE_2 t} \quad (2.30)$$

$$P_{ee} \equiv |\langle \nu_e | \nu_e \rangle|^2 = 1 - \sin^2 2\theta_V \sin^2 \left(\frac{1}{2}(E_1 - E_2)t \right) \quad (2.31)$$

where $E_1 - E_2 = (m_1^2 - m_2^2)/(2E) \equiv \pm \Delta m^2/(2E)$. In convenient units,

$$P_{ea} = \sin^2 2\theta_V \sin^2 \left(1.27 \frac{\Delta m^2 [\text{eV}^2]}{E_\nu [\text{GeV}]} L [\text{km}] \right). \quad (2.32)$$

2.2.2 Matter-Enhanced Oscillations

In the presence of matter, neutrinos oscillations are affected by scattering processes. Neutral current (NC) weak interactions with electrons and nucleons affect all flavors equally and introduce a physically-irrelevant phase, while ν_e alone may additionally participate in charged current (CC) interactions. The implications for flavor oscillations were first recognized by Wolfenstein [25] in 1978, but the idea did not gain traction until 1986 when Smirnov and Mikheyev realized that propagation in a medium with a varying electron density could lead to a resonant enhancement of flavor oscillations [26]. This effect is now known as the Mikheyev-Smirnov-Wolfenstein (MSW) effect, or simply the “matter effect.” A diagram representing the CC scattering responsible for the MSW effect is shown in Figure 2.1. To include the effects of matter, an effective potential is added to Hamiltonian given in Equation 2.29, such that $\mathcal{H}_{00} \rightarrow \mathcal{H}_{00} + 2\sqrt{2}G_F n_e E$, where G_F is the Fermi coupling constant (at low energies, the diagram in Figure 2.1 reduces to an effective quadrilinear coupling), n_e is the local electron density of the propagation medium, and E is the neutrino

³This assumption is valid in the case of solar neutrinos, but not in general. For a discussion of the generally-valid wave packet approach, see Reference [24].

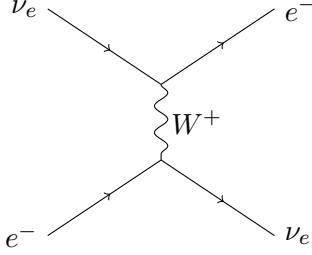


Figure 2.1: Charged-current $\nu_e - e^-$ scattering, which results in the resonant conversion of ν_e to other flavors in the MSW effect.

energy. Equation 2.29 may now be rewritten in terms of effective parameters [27]:

$$\mathcal{H} = \frac{\Delta \tilde{m}^2}{4E} \begin{pmatrix} -\cos 2\tilde{\theta} & \sin 2\tilde{\theta} \\ \sin 2\tilde{\theta} & \cos 2\tilde{\theta} \end{pmatrix}, \quad (2.33)$$

where

$$A = 2\sqrt{2}G_F n_e E \quad (2.34)$$

$$\cos \tilde{\theta} = \frac{-A/\Delta m^2 + \cos 2\theta_V}{\sqrt{(A/\Delta m^2 - \cos 2\theta_V)^2 + \sin^2 2\theta_V}} \quad (2.35)$$

$$\tilde{m}_1 = \frac{A}{2} - \frac{1}{2} \sqrt{(A/\Delta m^2 - \cos 2\theta_V)^2 + (\Delta m^2)^2 \sin^2 2\theta_V} \quad (2.36)$$

$$\tilde{m}_2 = \frac{A}{2} + \frac{1}{2} \sqrt{(A/\Delta m^2 - \cos 2\theta_V)^2 + (\Delta m^2)^2 \sin^2 2\theta_V}. \quad (2.37)$$

The resonance condition where

$$n_{e,res} = \frac{|\Delta m^2| \cos 2\theta_V}{2\sqrt{2}G_F E}, \quad (2.38)$$

2.2 Neutrino Mass and Mixing

implies that for certain neutrino energies and electron densities, large transition probabilities are possible even for small values of the vacuum mixing angle θ_V .

What happens to a neutrino traversing matter depends on values of the vacuum oscillation length L_V and the neutrino/electron interaction length L_e . If $L_V \ll L_e$, vacuum oscillations dominate; if $L_V \gg L_e$, mixing is suppressed; and if $L_V = L_e$, then

$$|\langle \nu_a | \nu_e \rangle|^2 = \sin^2(\pi R \sin 2\theta / L_V) \quad (2.39)$$

In regions of $n_e > n_{e,res}$, the extra potential means that ν_e is the heaviest eigenstate, and if the electron density decreases slowly along the neutrino path, the ν_e may be adiabatically transformed into the ν_2 eigenstate of the free (vacuum) Hamiltonian. In the case of solar neutrinos, this means that sufficiently high-energy ν_e emerge from the Sun effectively as ν_2 .

As a neutrino crosses the resonance, the probability of changing to the other adiabatic mass eigenstate may be computed by numerical methods but is analytically calculable for certain simple electron density profiles. For an exponential profile (approximating the Sun) [28],

$$P_{jump} = \frac{\exp\left[-2\pi r_0 \frac{\Delta m^2}{2E} \sin^2 \theta\right] - \exp\left[-2\pi r_0 \frac{\Delta m^2}{2E}\right]}{1 - \exp\left[-2\pi r_0 \frac{\Delta m^2}{2E}\right]} \quad (2.40)$$

where $r_0 \sim 0.1R_\odot$ is a characteristic length for the change of electron density. The ν_e survival probability is conveniently expressed in terms of this level crossing probability by Parke's formula:

$$P_{ee} = \frac{1}{2} + \left(\frac{1}{2} - P_{jump}\right) \cos 2\tilde{\theta} \cos 2\theta \quad (2.41)$$

2.3 Current Mixing Parameters

Results from a recent global fit for the mixing parameters are given by Gonzales-Garcia, Maltoni, and Schwetz in [29], and reproduced in Table 2.2. Propagating these uncertainties to the PMNS matrix, the following ranges are allowed at the 3σ confidence level [29]:

$$|U| = \begin{pmatrix} 0.801 \rightarrow 0.845 & 0.514 \rightarrow 0.580 & 0.137 \rightarrow 0.158 \\ 0.225 \rightarrow 0.517 & 0.441 \rightarrow 0.699 & 0.614 \rightarrow 0.793 \\ 0.246 \rightarrow 0.529 & 0.464 \rightarrow 0.713 & 0.590 \rightarrow 0.776 \end{pmatrix}. \quad (2.42)$$

For the present analysis, the parameters recommended by the Particle Data Group are used [30]:

$$\sin^2 \theta_{12} = 0.304_{-0.013}^{+0.014} \quad (2.43)$$

$$\Delta m_{21}^2 = (7.53 \pm 0.18) \times 10^{-5} \text{ eV}^2 \quad (2.44)$$

$$\sin^2 \theta_{23} = 0.514_{-0.056}^{+0.056} \text{ (NH)} \quad (2.45)$$

$$\sin^2 \theta_{23} = 0.511 \pm 0.55 \text{ (IH)} \quad (2.46)$$

$$\Delta m_{32}^2 = (2.49 \pm 0.06) \times 10^{-3} \text{ eV}^2 \text{ (NH)} \quad (2.47)$$

$$\Delta m_{32}^2 = (2.42 \pm 0.06) \times 10^{-3} \text{ eV}^2 \text{ (IH)} \quad (2.48)$$

$$\sin^2 \theta_{13} = (2.19 \pm 0.12) \times 10^{-2} \quad (2.49)$$

where NH and IH refer to the normal and inverted possibilities for the neutrino mass ordering. The CP-violating phase δ has not yet been measured.

2.3 Current Mixing Parameters

	Normal Ordering ($\Delta\chi^2 = 0.97$) bfp $\pm 1\sigma$	3 σ range	Inverted Ordering (best fit) bfp $\pm 1\sigma$	3 σ range	Any Ordering 3 σ range
$\sin^2 \theta_{12}$	$0.304^{+0.013}_{-0.012}$	$0.270 \rightarrow 0.344$	$0.304^{+0.013}_{-0.012}$	$0.270 \rightarrow 0.344$	$0.270 \rightarrow 0.344$
$\theta_{12}/^\circ$	$33.48^{+0.78}_{-0.75}$	$31.29 \rightarrow 35.91$	$33.48^{+0.78}_{-0.75}$	$31.29 \rightarrow 35.91$	$31.29 \rightarrow 35.91$
$\sin^2 \theta_{23}$	$0.452^{+0.052}_{-0.028}$	$0.382 \rightarrow 0.643$	$0.579^{+0.025}_{-0.037}$	$0.389 \rightarrow 0.644$	$0.385 \rightarrow 0.644$
$\theta_{23}/^\circ$	$42.3^{+3.0}_{-1.6}$	$38.2 \rightarrow 53.3$	$49.5^{+1.5}_{-2.2}$	$38.6 \rightarrow 53.3$	$38.3 \rightarrow 53.3$
$\sin^2 \theta_{13}$	$0.0218^{+0.0010}_{-0.0010}$	$0.0186 \rightarrow 0.0250$	$0.0219^{+0.0011}_{-0.0010}$	$0.0188 \rightarrow 0.0251$	$0.0188 \rightarrow 0.0251$
$\theta_{13}/^\circ$	$8.50^{+0.20}_{-0.21}$	$7.85 \rightarrow 9.10$	$8.51^{+0.20}_{-0.21}$	$7.87 \rightarrow 9.11$	$7.87 \rightarrow 9.11$
$\delta_{\text{CP}}/^\circ$	306^{+39}_{-70}	$0 \rightarrow 360$	254^{+63}_{-62}	$0 \rightarrow 360$	$0 \rightarrow 360$
$\frac{\Delta m_{21}^2}{10^{-5} \text{ eV}^2}$	$7.50^{+0.19}_{-0.17}$	$7.02 \rightarrow 8.09$	$7.50^{+0.19}_{-0.17}$	$7.02 \rightarrow 8.09$	$7.02 \rightarrow 8.09$
$\frac{\Delta m_{3\ell}^2}{10^{-3} \text{ eV}^2}$	$+2.457^{+0.047}_{-0.047}$	$+2.317 \rightarrow +2.607$	$-2.449^{+0.048}_{-0.047}$	$-2.590 \rightarrow -2.307$	$\begin{bmatrix} +2.325 \rightarrow +2.599 \\ -2.590 \rightarrow -2.307 \end{bmatrix}$

Table 2.2: Three-flavor oscillation parameters from our fit to global data after the NOW 2014 conference. The results are presented for the “Free Fluxes + RSBL” in which reactor fluxes have been left free in the fit and short baseline reactor data (RSBL) with $L \lesssim 100$ m are included. The numbers in the 1st (2nd) column are obtained assuming NO (IO), *i.e.*, relative to the respective local minimum, whereas in the 3rd column we minimize also with respect to the ordering. Note that $\Delta m_{3\ell}^2 \equiv \Delta m_{31}^2 > 0$ for NO and $\Delta m_{3\ell}^2 \equiv \Delta m_{32}^2 < 0$ for IO. Table and caption reproduced from [29]. “bfp” refers to the best-fit parameter value.

2.4 Status and Prospects

The tight constraints on the oscillation parameters given in the previous section demonstrate that in some respects the neutrino field is entering the regime of precision measurement: the mixing angles and squared mass splittings are known with impressive and improving precision. However, there remain fundamental questions in neutrino physics, for example:

Neutrino mass Why are the neutrinos so light in comparison the other Standard Model fermions? Is something other than the Higgs mechanism responsible for neutrino mass generation?

Neutrino Mass Ordering Is m_1 the lightest (normal mass ordering) or is m_3 (inverted mass ordering)?

Majorana neutrinos Is the neutrino a Majorana fermion (i.e. its own antiparticle), and if not, is lepton number a fundamental symmetry of the Standard Model?

CP Violation Could CP violation in neutrino interactions explain the matter-antimatter asymmetry in the universe by facilitating leptogenesis?

Sterile Neutrinos Are there additional neutrino types which participate in oscillations but not weak interactions, as some experimental evidence suggests?

NSI Are there weak nonstandard interactions in the neutrino sector that could be probed with long-baseline interferometry?

From its humble beginnings rescuing energy conservation in nuclear beta decay, the little neutrino has become a powerful tool for understanding fundamental physics on energy scales from μeV to the GUT scale, and from nuclear to cosmological spatial scales. Experiments under development today will measure the CP-violating phase δ_{CP} and determine the

2.4 Status and Prospects

ordering. Next-generation experiments may observe neutrinoless double-beta decay, demonstrating the Majorana nature of the neutrino and perhaps opening a window to GUT-scale physics generating the neutrino mass, at energies inconceivable for terrestrial colliders.

Chapter 3

Solar Neutrinos

The Sun-Earth system provides an incredible laboratory for precision low-energy neutrino oscillation measurements, with incredibly rich phenomenology. Electron neutrinos are produced in several of the nuclear reactions that power the Sun, and propagate through a nonuniform electron-rich medium, undergoing flavor- and energy-dependent scattering which can resonantly enhance the flavor oscillations. The neutrinos are produced at a wide range of energies and radii, making the resultant flux on Earth sensitive to details of the solar composition and structure.

This modern perspective was hard-won over decades of experiments; the deficit in flux now attributed to flavor oscillations, known as the “solar neutrino problem,” stood unresolved from the first flux observations in the late 1960s until the Sudbury Neutrino Observatory results in 2001 [6].

3.1 Standard Solar Models

Solar models predict the evolutionary history of the Sun, assuming that it is driven by well-understood nuclear physics and fluid dynamics. The models take as input experimentally-determined quantities such as nuclear cross sections and a set of parameters like the initial abundances of light elements in the Sun, and use the observational boundary conditions (photon luminosity, size, age) to find the most suitable parameters. These parameters then predict the neutrino fluxes, in terms of both energy spectra and where the reactions occur within the volume of the sun.

The cornerstone of these models is the postulate that the Sun’s energy comes from a proton-proton (*pp*) fusion chain terminating with the production of an α particle. There are several branches in the *pp* chain in which reactions produce ν_e with a range of energies up to ~ 18 MeV. A second chain which fuses protons into α particles through catalysis by heavier elements is the Carbon-Nitrogen-Oxygen (CNO) cycle. This set of reactions is subdominant in the Sun but is the primary mechanism for energy production in slightly larger stars. In either case, the overall process is that of “hydrogen burning:”



The *pp* chain reactions are shown in Figure 3.1 and the neutrino energy spectra for the *pp* and CNO cycles are in Figure 3.2. The spatial distribution of the fluxes is shown in Figure 3.3 for the GS98SFII model⁴. Table 3.1 summarizes the predicted ν_e fluxes for each neutrino-producing reaction in the *pp* and CNO chains, for selected SSMs. For the present analysis, The BSB05(GS98) – henceforth denoted BS05(OP) – is used unless noted otherwise.

⁴This model comprises the GS98 model proposed by Grevesse and Sauval in 1998 [31], with the updated nuclear reaction rate calculations known as Solar Flux II, or SFII [32].

3.1 Standard Solar Models

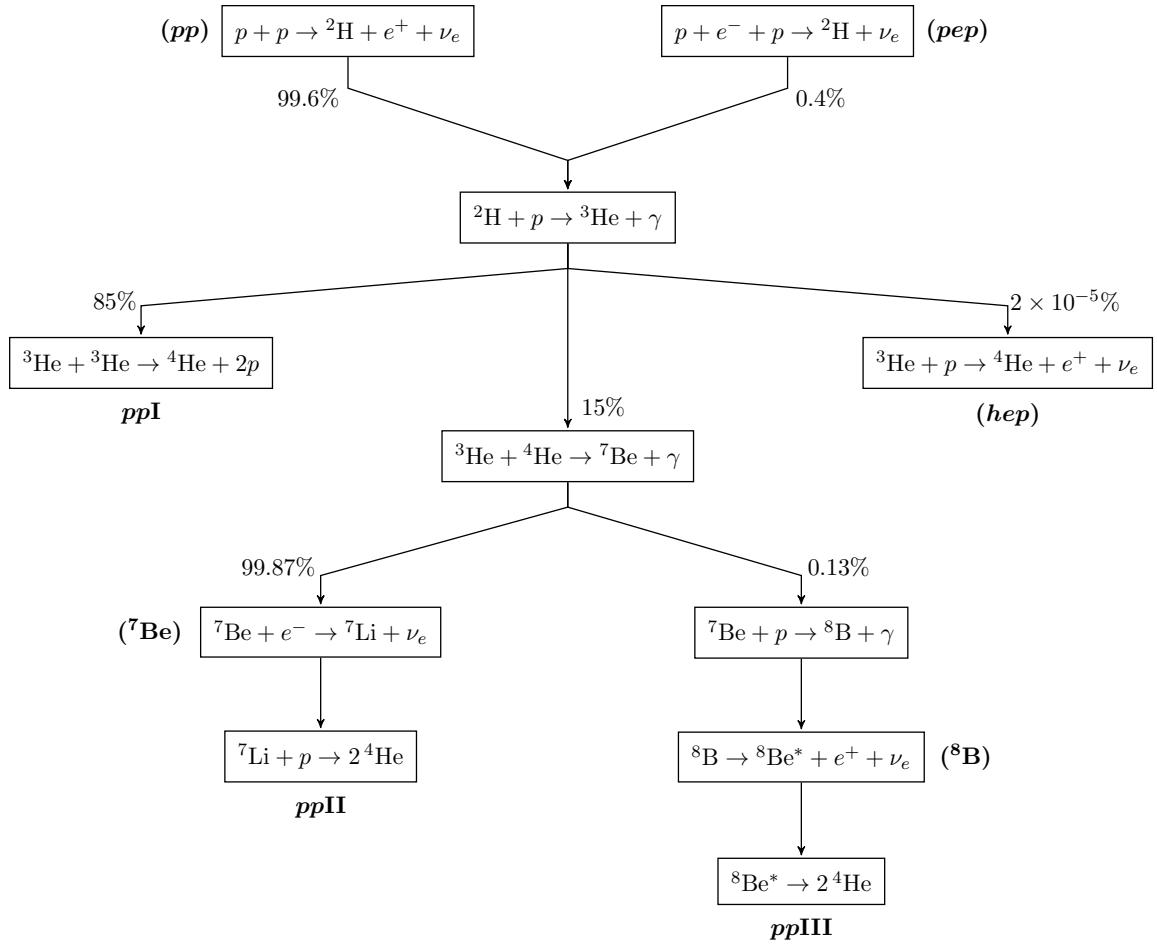


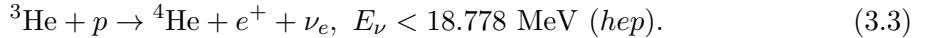
Figure 3.1: The reactions of the solar *pp* fusion chain, with the names of the neutrino-producing reactions are given in parentheses. Figure based on Fig. 10.1 of Reference [21].

3.2 Solar Neutrino Oscillations

	BP00 [34]	BP04 [35]	BSB05(GS98) [36]	BSB05(AGS05) [36]
$\Phi_{pp} / 10^{10}$	$5.95(1 \pm 0.01)$	$5.94(1 \pm 0.01)$	$5.99(1 \pm 0.009)$	$6.06(1 \pm 0.007)$
$\Phi_{pep} / 10^8$	$1.40(1 \pm 0.015)$	$1.40(1 \pm 0.02)$	$1.42(1 \pm 0.015)$	$1.45(1 \pm 0.012)$
$\Phi_{hep} / 10^3$	9.3	$7.88(1 \pm 0.16)$	$7.93(1 \pm 0.155)$	$8.25(1 \pm 0.155)$
$\Phi_{7\text{Be}} / 10^9$	$4.77(1 \pm 0.10)$	$4.86(1 \pm 0.12)$	$4.84(1 \pm 0.105)$	$4.34(1 \pm 0.093)$
$\Phi_{8\text{B}} / 10^6$	$5.05(1^{+0.20}_{-0.16})$	$5.79(1 \pm 0.23)$	$5.69(1^{+0.173}_{-0.147})$	$4.51(1^{+0.127}_{-0.113})$
$\Phi_{13\text{N}} / 10^8$	$5.48(1^{+0.21}_{-0.17})$	$5.71(1^{+0.37}_{-0.35})$	$3.05(1^{+0.366}_{-0.268})$	$2.00(1^{+0.145}_{-0.127})$
$\Phi_{15\text{O}} / 10^8$	$4.80(1^{+0.25}_{-0.19})$	$5.03(1^{+0.43}_{-0.39})$	$2.31(1^{+0.374}_{-0.272})$	$1.44(1^{+0.165}_{-0.142})$
$\Phi_{17\text{F}} / 10^6$	$5.63(1 \pm 0.25)$	$5.91(1 \pm 0.44)$	$5.83(1^{+0.724}_{-0.420})$	$3.25(1^{+0.166}_{-0.142})$

Table 3.1: Solar neutrino fluxes in $\text{cm}^{-2} \text{s}^{-1}$ for a variety of SSMs. Table reproduced from Reference [21].

The two reactions of primary interest for the present analysis are



The *hep* reaction produces the highest-energy solar neutrinos, but with a very small branching ratio (2×10^{-7} per *pp* termination). This will produce an excess of detected events above the endpoint of ${}^8\text{B}$ spectrum, but with a flux about three orders of magnitude smaller.

3.2 Solar Neutrino Oscillations

The *pp* and CNO reactions form a pure ν_e source, with each reaction having a characteristic neutrino energy spectrum and radial distribution within the stellar interior, and these solar neutrinos initially in the ν_e flavor eigenstate undergo flavor oscillations en route to Earth.

3.2 Solar Neutrino Oscillations

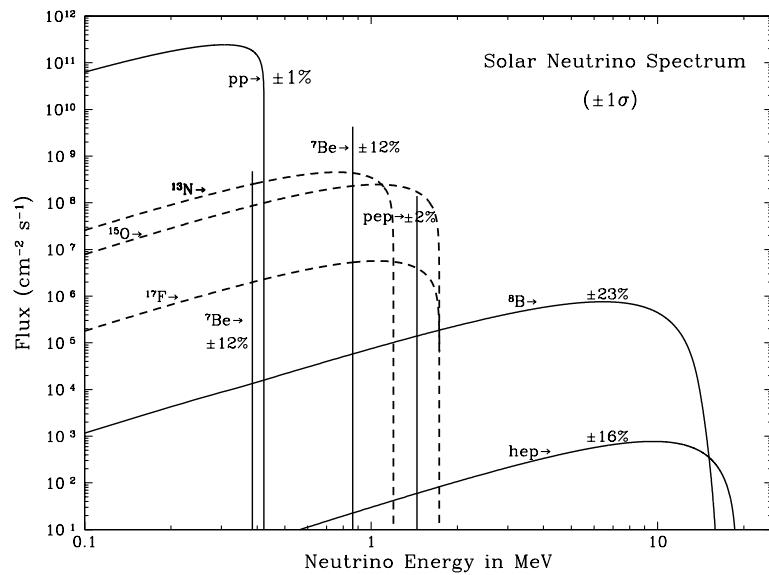


Figure 3.2: Solar neutrino spectrum. This figure shows the energy spectrum of neutrinos predicted by the standard solar model (Bahcall and Pinsonneault 2004). The neutrino fluxes from continuum sources (like pp and 8B) are given in the units of number per cm^2 per second per MeV at one AU. The line fluxes (pep and ^7Be) are given in number per cm^2 per second. The spectra from the pp chain are drawn with solid lines; the neutrino energy spectra from reactions with carbon, nitrogen, and oxygen (CNO) isotopes are drawn with dotted lines. Figure from Reference [33].

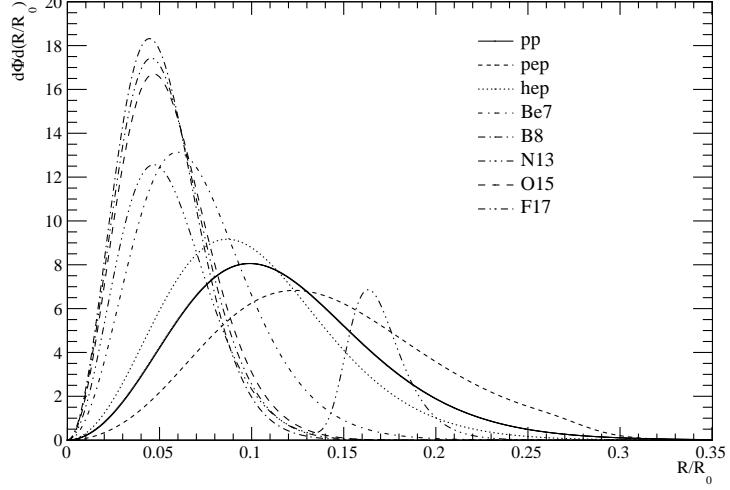


Figure 3.3: Spatial distributions of the fluxes in the GS98SFII standard solar model.

For real systems including the Sun-Earth system, the survival probability must be averaged over the finite energy distribution and spatial extent of the source. The oscillating term in Equation 2.32 are replaced with an integral over the solar model [21]:

$$P_{ea} = \frac{1}{2} \sin^2 2\theta \left[1 - \left\langle \cos \left(\frac{\Delta m^2 L}{2E} \right) \right\rangle \right] \quad (\alpha \neq \beta) \quad (3.4)$$

where

$$\left\langle \cos \left(\frac{\Delta m^2 L}{2E} \right) \right\rangle = \int \cos \left(\frac{\Delta m^2 L}{2E} \right) \phi(L/E) d(L/E). \quad (3.5)$$

Averaging over the broad energy and spatial distributions for the ${}^8\text{B}$ and *hep* fluxes, the ν_e survival probability for vacuum oscillations becomes

$$\langle P_{ee} \rangle = 1 - \frac{1}{2} \sin^2 2\theta_{12} \quad (3.6)$$

3.3 The Homestake Experiment

in the effective two-flavor theory, or

$$\langle P_{ee}^{(3\nu)} \rangle = \sin^4 \theta_{13} + \cos^4 \theta_{13} \left(1 - \frac{1}{2} \sin^2 2\theta_{12} \right) \quad (3.7)$$

in the full three-neutrino scenario.

Due to the energy dependence of the MSW effect described in Section 2.2.2, neutrinos created at lower energies emerge from the Sun effectively as ν_e and will undergo such vacuum oscillations, while neutrinos created at higher energies are effectively ν_2 , which is an admixture with substantial fraction of ν_e , ν_μ , and ν_τ , about 1/3 each. The survival probability in this limit is

$$\langle P_{ee} \rangle = \cos^4 \theta_{13} \sin^2 \theta_{12}. \quad (3.8)$$

Between these limits, the energy dependence of the survival probability may be determined numerically, by integrating along neutrino paths considering the spatial and energy distributions and the electron density profile for a given solar model. Figure 3.4 shows for example P_{ee} for ${}^8\text{B}$ solar neutrinos assuming the BS05(OP) SSM, indicating the vacuum-dominated and matter-dominated limits. The transition region is highly sensitive to the effective potential for solar ν_e , making it an interesting probe for non-standard neutrino interactions (NSI) as well as non-standard solar physics. NSI models are poorly constrained by current data, according to a recent global analysis [37].

3.3 The Homestake Experiment

Beginning in 1962, John Bahcall and collaborators had developed robust SSMs making definite predictions of the solar neutrino flux on Earth. In the late 1960s, an experiment

3.3 The Homestake Experiment

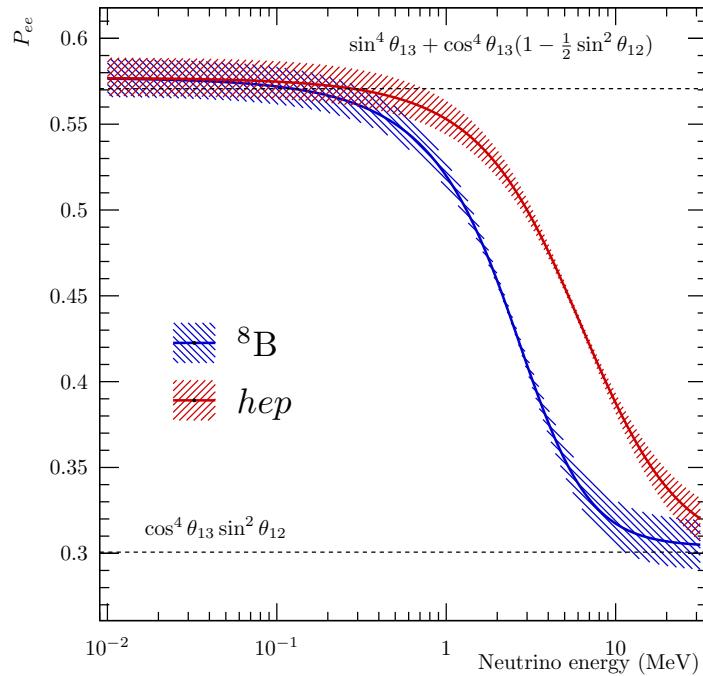
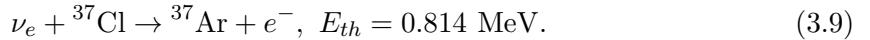


Figure 3.4: The survival probability P_{ee} for ${}^8\text{B}$ and hep solar neutrinos, assuming the BS05(OP) SSM, and PDG-recommended three-neutrino mixing parameters (with their uncertainties shown as bands), according to the Sun-Earth Large Mixing Angle Adiabatic (SELMAA) calculation (see Reference [38]).

3.4 The Solar Neutrino Problem

was devised by Raymond Davis and collaborators that would detect this flux in order to make a detailed measurement of stellar thermonuclear processes.⁵

The Homestake experiment [1] was located at the Homestake gold mine near Lead, SD, at a depth of 4200 meters water equivalent (m.w.e.), and operated from 1970–1994. Solar ν_e were observed through the production of ^{37}Ar by a charged-current interaction of ν_e on ^{37}Cl :



The argon atoms were extracted periodically from the target vessel and gaseous samples were assayed in miniature proportional counters. Samples were taken to be counted periodically, with 108 individual measurements occurring throughout the duration of the experiment. Reference [1] reports best-fit production rates for both a fit to the full data set and for each individual run. It soon became clear that the Homestake experiment saw a deficit in the detected flux, by about a factor of three. This discrepancy persisted through systematic checks of the Ar extraction efficiency and analysis improvements (pulse-shape discrimination for better SNR in the proportional counters), and thus was born the Solar Neutrino Problem.

3.4 The Solar Neutrino Problem

The deficit in the observed solar neutrino flux known as the Solar Neutrino Problem persisted for decades, motivating several major experiments before SNO definitively identified matter-enhanced flavor oscillations as the cause; solar ν_e changed to ν_μ and ν_τ invisible to previous detectors. Naturally, alternative hypotheses were proposed in the mean time, many of which attempted to preserve the massless Standard Model neutrino. A few notable classes of theories included:

⁵This followed a similar experiment at Savannah River searching for reactor neutrinos with ^{37}Cl , which helped establish that ν and $\bar{\nu}$ interact differently [39].

3.4 The Solar Neutrino Problem

Nonstandard Solar Models The Standard Solar Model (SSM) was known to be imperfect (and remains so). Even in the 1970s, varying the model parameters conspiratorially to suppress the neutrino flux was unable to account for the discrepancy. However, modifications were proposed that suppressed the neutrino fluxes while respecting the observational boundary conditions, for example disequilibrium in the chemical composition of the Sun. Such models were generally ruled out by improved observational constraints.

Spin-flip and Neutrino Decay Two more exotic exotic explanations for a low ν_e count rate are spin-flip effects and neutrino decay. The former, of which there is a standard and matter-enhanced variety, relies on a large neutrino magnetic moment: if this exists, a solar ν_e could change helicity en route to Earth and thus become undetectable by the Cl and Ga experiments. The latter, neutrino decay, could explain the deficit if the neutrino lifetime were sufficiently short that neutrinos of the heavier mass state decayed into a lighter antineutrino and some scalar boson. The data at the time did not exclude this explanation, but constrained it to a small and unexpectedly high range of mixing angles.

WIMPs Weakly Interacting Massive Particles, or WIMPs, were proposed at as a simultaneous solution to both the Solar Neutrino Problem and the Dark Matter Problem. The argument was that energy transport by WIMPs, which was not accounted for in any solar models, could reduce the thermal gradient in the Sun sufficiently to lower the ${}^8\text{B}$ neutrino flux to the observed value.

It was not until the mid-1980s that the MSW theory was generally accepted as a plausible solution to the solar neutrino problem. Assuming vacuum oscillations only, Equation 3.6 gives the average ν_e survival probability. Using a modern value of $\theta_V \sim 33^\circ$, $\langle P_{ee} \rangle \sim 0.6$, which does not explain the factor of ~ 3 suppression observed at Homestake. If instead the Cabibbo angle, the known mixing angle at the time the problem arose, is assumed,

$\theta_V \sim \theta_C \sim 13^\circ$ and $\langle P_{ee} \rangle \sim 0.9$, which appears not very promising at all. Furthermore, the initial papers on the MSW effect contained a sign error, such that resonant flavor conversion only occurred for the unexpected case where $m(\nu_e) > m(\nu_\mu), m(\nu_\tau)$; this was corrected by Langacker in 1986 [12].

3.5 Experimental Results

A number of other solar neutrino detectors followed Homestake, aiming to confirm or refute the anomalous rate measurement, and it was also recognized that water Cherenkov detectors built to search for nucleon decay could be repurposed to make real-time solar neutrino measurements. This section briefly summarizes major results in the field; a list is given in Table 3.2.

The next solar neutrino experiments to come online after Homestake — Kamiokande II [40], SAGE [4, 41] and GALLEX/GNO [2, 3, 42] — confirmed the existence of the Solar Neutrino Problem. Super-Kamiokande established strong evidence for neutrino oscillations through observation of atmospheric neutrinos [19], and SNO demonstrated that the SNP was caused by resonantly-enhanced oscillations of solar neutrinos [6]. Following that discovery, SNO [43–48] and Super-Kamiokande [5, 49–52] have made precision measurements of solar neutrino oscillations, and the scintillator-based Borexino experiment has pushed down energy thresholds to make real-time measurements of the ${}^7\text{Be}$ [53], ${}^{13}\text{C}$ [54], and very low-energy pp [55] neutrinos. A number of next-generation experiments, including SNO+ [56] and THEIA [57], have potential to greatly improve the precision in the coming years.

3.5.1 Searches for *hep* Neutrinos

Searches for *hep* neutrinos also have an interesting history. In March 2000, with about 800 days of data, the Super-Kamiokande collaboration reported an excess in the high-energy

3.5 Experimental Results

Name	Dates	Target	Mass	Location
Homestake	1970–1994	^{37}Cl	615 t C_2Cl_4	Lead, SD
SAGE	1989–2007	^{71}Ga	60 t GaCl_3	Baksan
GALLEX	1991–1997	^{71}Ga	30.3 t $\text{GaCl}_3\text{--HCl}$	LNGS
GNO	1998–2003	^{71}Ga	30.3 t $\text{GaCl}_3\text{--HCl}$	LNGS
Kamiokande II	1987–1996	H_2O	2.1 kt	Kamioka
Super-Kamiokande	1996–	H_2O	50 kt	Kamioka
SNO	2001–2006	D_2O	1 kt	Sudbury
Borexino	2007–	LS	278 t	LNGS
SNO+ (Future)	2017–	LS	780 t	Sudbury

Table 3.2: Summary of solar neutrino flux measurements.

spectrum, a low-significance hint of an unexpectedly high flux of *hep* solar neutrinos, at the level of 16.7 times the (BP98) SSM [58]. The fit and energy spectrum are shown in Figure 3.5.

The first SK publication on the subject came in 2001, with 1258 days of data, and the discrepancy apparently resolved. A limit of $4.0 \times 10^4 \nu \text{ cm}^{-2} \text{ s}^{-1}$ (90% CL), or 4.3 times the BP2000 SSM prediction, was obtained for the *hep* flux [60] based on data shown in Figure 3.6. In 2006, SK published a revised limit based on the full SK-I data set. With 4.9 ± 2.7 events in the energy range from 18–21 MeV, the 90% CL limit was slightly weakened, at $7.3 \times 10^4 \nu \text{ cm}^{-2} \text{ s}^{-1}$ [50].

Also in 2006, the SNO collaboration published a stronger limit based on an exposure of 0.65 kilotonne-years of a pure D_2O target, $2.3 \times 10^4 \nu \text{ cm}^{-2} \text{ s}^{-1}$ at the 90% CL [15], with two events observed in a 14.3–20 MeV signal region. The data and Monte Carlo prediction are shown in Figure 3.7. This dissertation expands on that result, as detailed in the subsequent chapters.

3.5 Experimental Results

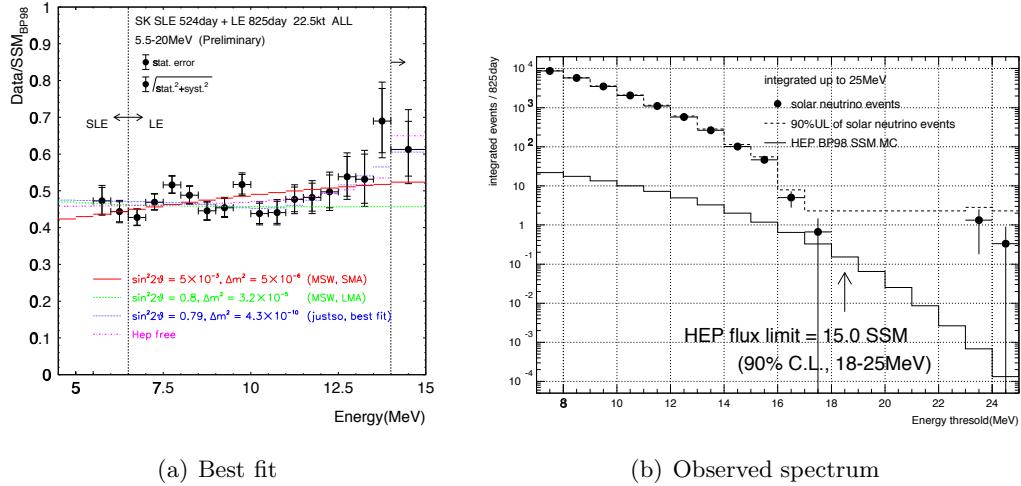


Figure 3.5: The initial Super-Kamiokande data showing a potential *hep* excess. Figures from Reference [59].

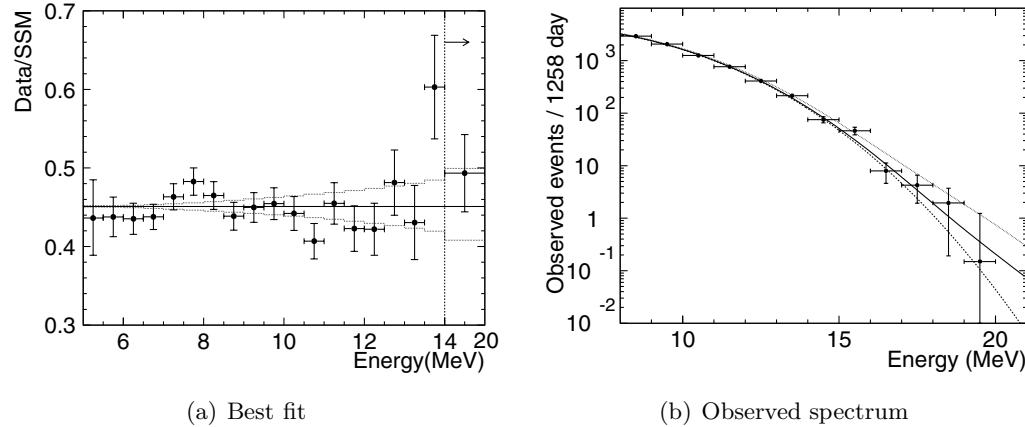


Figure 3.6: The published 2001 Super-Kamiokande data showing a *hep* flux consistent with the SSM. Figures from Reference [60].

3.5 Experimental Results

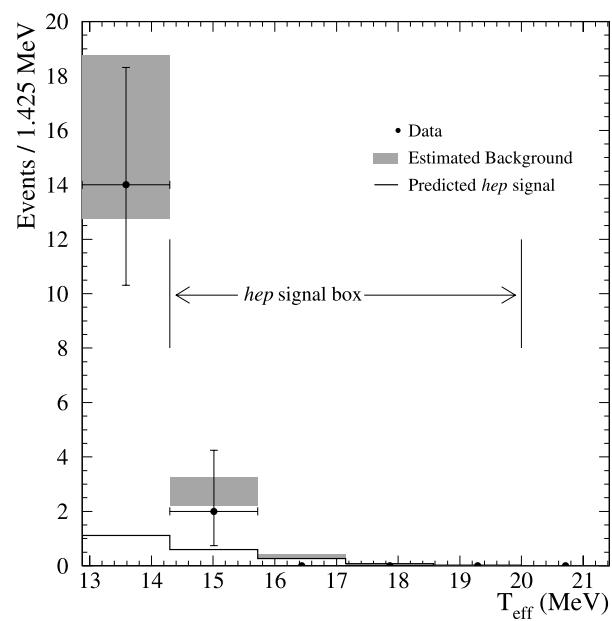


Figure 3.7: The spectrum published by SNO for events near the ${}^8\text{B}$ endpoint, indicating the hep search region. Shaded bands show the systematic uncertainties. Figure from Reference [15].

Chapter 4

The Diffuse Supernova Neutrino Background

A core-collapse supernova (SN) is a dramatic astrophysical event wherein a star, no longer supported by thermal or electron degeneracy pressure, collapses to a neutron star or black hole, ejecting its outer layers. While supernovae can be so optically bright as to outshine their host galaxies, most of the energy loss occurs through neutrinos, which interact only weakly with the surrounding medium. Nearby core-collapse supernovae — those close enough that terrestrial neutrino detectors would observe a statistically significant number of interactions — are rare events, with only a few expected per century. The last such event, Supernova 1987A, provided a wealth of information with only ~ 20 events detected worldwide, but opened new questions as well [13].

The diffuse supernova neutrino background⁶ (DSNB) is the “glow” of neutrinos emitted

⁶This is occasionally referred to as the *relic* neutrino background, or the relic supernova neutrino back-

in the large number of supernovae in the visible universe that are not detectable on an event-by-event basis. Models generally suggest that this signal should be at the limits of detection of current-generation neutrino observatories, and likely accessible to the next generation. A measurement of the DSNB flux would provide valuable input for modeling of core-collapse supernovae: the spectrum carries information about the mean temperature and average total energy of these systems, and the flavor composition could provide insight into the weak processes involved and help to constrain models of supernova dynamics.

As the DSNB has not yet been observed, we must defer the spectroscopy and flavor characterization to a future in which more sensitive detectors than SNO observe a significant sample of DSNB events. The details of supernova dynamics, and the flavor-changing processes that occur within, are at this time very uncertain. I therefore choose for this analysis a simple but illustrative model, and report counting rates in the relevant energy region, to will allow freedom of interpretation once the source effects are clarified.

4.1 Core-Collapse Supernovae

The structure and size of a star is determined by the balance of thermal, electron degeneracy, and gravitational pressure. Throughout the stellar life cycle, the core will fuse heavier and heavier elements; once a source is exhausted, the core contracts, heats up, and fusion of the next becomes energetically available. In sufficiently massive stars, this proceeds until the chain reaches iron, which has the highest binding energy. At this point, the fate of the star depends on its mass, which determines whether the gravitational pressure can overcome the electron degeneracy pressure, leading to a collapse to either neutron star or black hole. The

ground (RSNB). This must not be confused with the cosmological relic neutrinos which are a product of the Big Bang.

4.1 Core-Collapse Supernovae

threshold mass is known as the Chandrasekhar limit, and given by [21]

$$M_C \simeq 5.83 Y_e^2 M_\odot \simeq 1.46 M_\odot \quad (4.1)$$

where $Y_e = N_p/(N_p + N_n) \sim 0.5$ is the electron fraction and M_\odot is the mass of the Sun (one solar mass). In the case of core-collapse supernovae, the instability occurs once this mass of iron has been produced [61]. The specific type of supernova, characterized by the light curve and emission spectrum, as well as the type of remnant (neutron star or black hole) depends on the initial mass and the composition (specifically the metallicity) of the progenitor star.

The rebound following the collapse triggers an outward-directed shock wave, which dissociates the material in the collapsing outer layers. Electron neutrinos produced in electron captures on the free protons behind the high-density shock front are trapped, emerging as a burst — the neutronization or shock breakout burst — once the expansion reaches a point where the density becomes low enough. This “prompt” burst is a relatively small component of the total neutrino emission, although the timing is very informative in single-event observations. The bulk of supernova neutrinos are produced in the collapsed core via numerous $\nu\bar{\nu}$ pair-producing reactions, leading to roughly equal numbers of each flavor. Due to the high (nuclear) density of this environment, neutrinos emitted in the core will diffuse outward, undergoing many scatters, and therefore emerge from the surface of last scattering (known as the neutrinosphere) with an approximately thermal spectrum [61]:

$$\frac{dN(E)}{dE} = \frac{E_\nu^{\text{tot}}}{6} \frac{120}{7\pi^4} \frac{E^2}{T^4} \left[\exp\left(\frac{E}{T}\right) + 1 \right]^{-1}, \quad (4.2)$$

where E_ν^{tot} is the total energy emitted in neutrinos (assuming equipartition of flavors) and T is the temperature at the surface of last scattering for the neutrino flavor under consid-

eration. Detailed calculations suggest that the spectrum is slightly squeezed (suppressed at low and high temperatures), however the data required to constrain such models are sparse [62].

4.2 The Diffuse Supernova Neutrino Background

Production The DSNB flux is calculated with a line-of-sight integral of sources, accounting for the variation in the intrinsic rate of supernovae as a function of redshift. The flux expected on Earth is given by [63]

$$\frac{d\phi(E)}{dE} = \int R_{SN}(z) \frac{dN[E(1+z)]}{dE} (1+z) \left| \frac{dt}{dz} \right| dz \quad (4.3)$$

where $R_{SN}(z)$ is the intrinsic rate of core-collapse supernovae at redshift z , dN/dE is the (thermal) emission spectrum, and $|dt/dz|$ depends on cosmological parameters. In a minimal Λ CDM model,

$$\left| \frac{dt}{dz} \right| = \left[H_0(1+z) \sqrt{\Omega_M(1+z)^3 + \Omega_\Lambda} \right]^{-1}, \quad (4.4)$$

with H_0 the Hubble constant and Ω_M and Ω_Λ the relative energy density due to matter and the dark energy, respectively. These parameters have been measured by the Planck experiment [64], with results summarized in Table 4.1.

In the “concordance” model of L. Strigari [65], the intrinsic rate of core-collapse supernovae, R_{SN} , is inferred from measurements of the UV, visible, and IR luminosity of galaxies. The formation rate of massive stars is measured through these channels, and it is assumed that the death rate is equal such that the galaxy remains in equilibrium. Then, an initial mass function (IMF) is used as a model of the population of stellar masses for extrapolation

4.2 The Diffuse Supernova Neutrino Background

Parameter	TT+lowP	TT,TE,EE+lowP+lensing+ext
H_0	67.31 ± 0.96	67.74 ± 0.46
Ω_M	0.315 ± 0.013	0.3089 ± 0.0062
Ω_Λ	0.685 ± 0.013	0.6911 ± 0.0062

Table 4.1: Cosmological parameters as measured by Planck in [64]. H_0 is given in units of $\text{km s}^{-1} \text{Mpc}^{-1}$. See reference for further details on the different constraints.

to lower-mass stars (down to the Chandrasekhar limit). A star formation rate of the form

$$R_{SF}(z) = \begin{cases} R_{SF}(0)(1+z)^\beta, & z < z_p \\ R_{SF}(z_p)(1+z)^\alpha, & z > z_p \end{cases} \quad (4.5)$$

with $\beta \sim 2$, $z_p \sim 2$, and $\alpha \sim 0$ — that is, the star formation rate is a power law up to some threshold redshift, after which it is constant — is found to be consistent with observations; see Figure 4.1.

Given R_{SF} and an IMF Ψ , the supernova rate is given by [61]

$$R_{SN}(z) = R_{SF} \frac{\int_8^{50} \Psi(M) dM}{\int_{0.1}^{100} M \Psi(M) dM} \simeq \frac{R_{SF}(x)}{143 M_\odot} \quad (4.6)$$

where the rightmost expression is calculated assuming the canonical Salpeter IMF,

$$\Psi(M) = \frac{dn}{dM} \propto M^{-2.35}, \quad 0.1 M_\odot < M < 100 M_\odot.$$

Putting this together with the approximate neutrino energy spectrum given in Equation 4.2 and computing the integral over redshift, we arrive at the DSNB flux prediction on Earth shown in Figure 4.2. As redshift increases, the spectrum shifts toward lower energy,

4.2 The Diffuse Supernova Neutrino Background

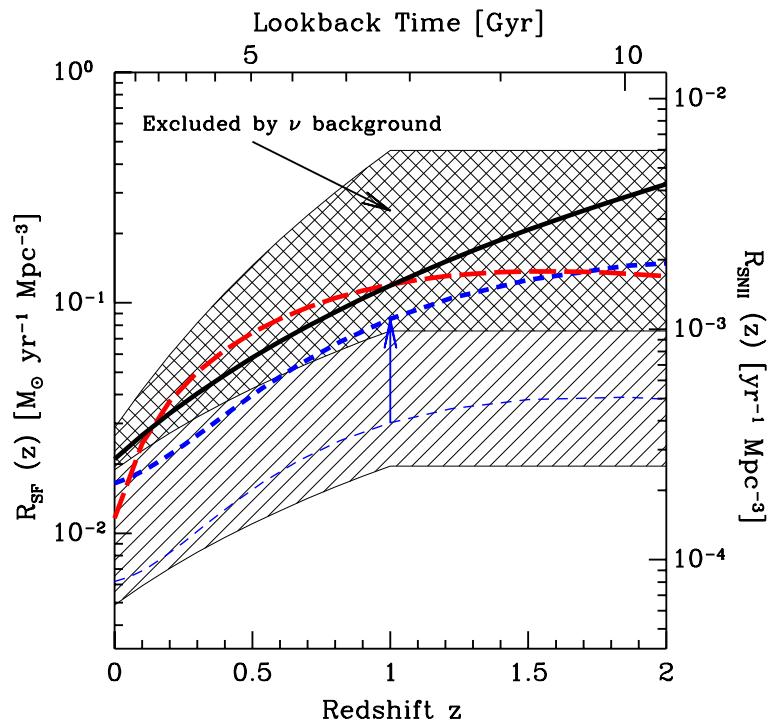


Figure 4.1: Comparison of data with the Strigari concordance model for the star formation rate. The hatched band is allowed by observations including the SuperK DSNB $\bar{\nu}_e$ limit. The black, red, and blue lines are measurements by Dahlen et al. [66], GALEX [67], and 2MASS+2dF [68], respectively, and indicating the effect of a dust correction for the latter. Figure from [65].

4.3 Experimental Efforts

and the contribution with $E_\nu > 15$ MeV is negligible for $z \gtrsim 3$, so the integral is cut off at that point. This cutoff affects the shape at low energy, where the DSNB is already a negligible background to the ^8B solar neutrinos.

An alternative approach, using the catalog of $z \lesssim 0.05$ Type II core-collapse supernovae (SNII) which were optically observed while SNO was taking data to perform a time-correlated search, is explored in Appendix C. The best sensitivity, however, is obtained using the method described above, due in general to the much larger volume of space being sampled. This technique may be of interest, however, for more sensitive searches such as at the Super-Kamiokande $\bar{\nu}_e$ search.

Detection The rate of DSNB ν_e interactions expected in a terrestrial detector is given by the product of the flux with the detector material cross sections, scaled by the number of targets [63]:

$$R(E) = \sum_T N_T t \int \frac{d\phi(E_\nu)}{dE_\nu} \frac{d\sigma_T(E_\nu, E_e)}{dE_e} dE_\nu dE_e \quad (4.7)$$

where T represents a type of target, t the exposure time, E_ν the incoming neutrino energy, and E_e the energy of an outgoing electron.

4.3 Experimental Efforts

Detection of the DSNB would not only provide a means of studying supernova neutrinos while waiting for another SN1987A-like event, but also provide crucial context to single-event observations. Therefore a number of experimental attempts have been made to observe this flux, and experimental sensitivities will meet or exceed the model predictions in the near future. As we have learned from the case of solar neutrinos, detecting the flux is

4.3 Experimental Efforts

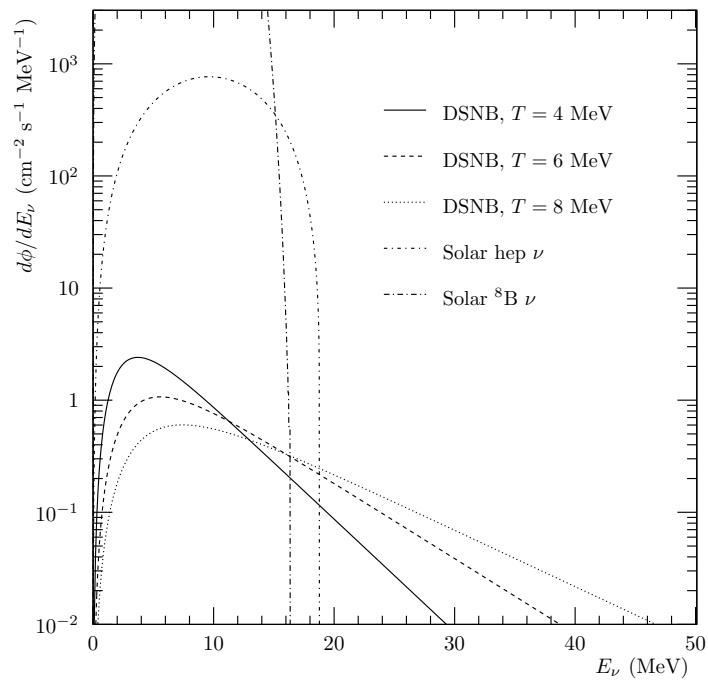


Figure 4.2: Expected flux of DSNB neutrinos on Earth, assuming a thermal spectrum. The unoscillated solar ${}^8\text{B}$ and hep fluxes are shown for scale. Compare to Figure 1 in [63].

4.3 Experimental Efforts

but the first step in understanding the source. This section reviews existing experimental searches for the DSNB and highlights some significant future prospects.

4.3.1 Current Results

The first direct limits on the DSNB fluxes came from the Liquid Scintillator Detector (LSD), a 90 ton liquid scintillator detector operated at Mont Blanc in the late 1980s to search for antineutrino bursts [69]. Search channels included inverse beta decay (IBD) for $\bar{\nu}_e$, charged-current interactions on ^{13}C for ν_e and $\bar{\nu}_e$, and neutral-current interactions with ^{12}C for $\nu_{\mu\tau}$ and $\bar{\nu}_{\mu\tau}$, all using time coincidences for event tagging. The LSD experiment reported the following flux limits in 1992 [69]:

$$\Phi_{\nu_e} < 6.8 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1}, \quad 25 < E_\nu < 50 \text{ MeV} \quad (4.8)$$

$$\Phi_{\bar{\nu}_e} < 9.0 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}, \quad 9 < E_\nu < 50 \text{ MeV} \quad (4.9)$$

$$\Phi_{\bar{\nu}_e} < 8.2 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1}, \quad 20 < E_\nu < 50 \text{ MeV} \quad (4.10)$$

$$\Phi_{\nu_{\mu\tau}} < 3.0 \times 10^7 \text{ cm}^{-2} \text{ s}^{-1}, \quad 20 < E_\nu < 100 \text{ MeV} \quad (4.11)$$

$$\Phi_{\bar{\nu}_{\mu\tau}} < 3.3 \times 10^7 \text{ cm}^{-2} \text{ s}^{-1}, \quad 20 < E_\nu < 100 \text{ MeV} \quad (4.12)$$

In 2006, in a combined search for *hep* solar neutrinos, the SNO experiment set an improved direct limit on the ν_e component of the flux [15]:

$$\Phi_{\nu_e} < 70 \text{ cm}^{-2} \text{ s}^{-1}, \quad 22.9 < E_\nu < 36.9 \text{ MeV} \quad (4.13)$$

This remains the strongest direct ν_e constraint to date, and that work is the foundation for the DSNB search presented in this dissertation.

4.3 Experimental Efforts

The strongest overall limits on the DSNB flux are due to the Super-Kamiokande (SK) experiment, based on a search for $\bar{\nu}_e$ through IBD. While SK places direct limits on the $\bar{\nu}_e$ flux, this limit has also been used to derive model-dependent limits on the ν_e [70], $\nu_{\mu\tau}$, and $\bar{\nu}_{\mu\tau}$ fluxes in the range $E_\nu > 19.3$ MeV [71]:

$$\Phi_{\nu_e} < 73.3 - 154 \text{ cm}^{-2} \text{ s}^{-1} \quad (4.14)$$

$$\Phi_{\bar{\nu}_e} < 1.4 - 1.9 \text{ cm}^{-2} \text{ s}^{-1} \quad (4.15)$$

$$\Phi_{\nu_{\mu\tau}} < (1.0 - 1.4) \times 10^3 \text{ cm}^{-2} \text{ s}^{-1} \quad (4.16)$$

$$\Phi_{\bar{\nu}_{\mu\tau}} < (1.3 - 1.9) \times 10^3 \text{ cm}^{-2} \text{ s}^{-1} \quad (4.17)$$

where the limits represent bounds at the 90% confidence level. In the range $22.9 < E_\nu < 36.9$, the ν_e flux limit is $\Phi_{\nu_e} < 39 - 54 \text{ cm}^{-2} \text{ s}^{-1}$.

4.3.2 Future Directions

Recently, the community's attention has focused on a Gd-loaded phase of SK (SK-Gd) [72]. This will improve the sensitivity by enhancing the neutron detection efficiency, and given the proximity of the limit to the theoretical prediction, there is well-founded hope that this will bring the first observation of the DSNB. Following a significant R&D effort, SK-Gd has been approved to proceed as of June 2015 [73].

While a tremendous achievement, this measurement would only directly address the $\bar{\nu}_e$ component of the flux, leaving much work to do for a complete picture of neutrino production and propagation in supernovae. Next-generation detectors such as DUNE (LAr target) [74] and THEIA (water-based liquid scintillator with potential for target isotope loading) [57] could help to improve the precision of the flux measurement and begin to address the question of flavor composition. With SK-Gd moving forward, large LAr TPC-based de-

4.3 Experimental Efforts

tectors slated to begin operation in the near future, and LS detectors of unprecedented scale on the horizon, prospects are good for not only detection, but also quantitative spectral measurements that will constrain SN models in the coming years.

Chapter 5

The Sudbury Neutrino Observatory

The Sudbury Neutrino Observatory experiment, originally envisioned by H. Chen in 1985 [75], was built to conclusively solve the solar neutrino problem outlined in Chapter 3. Whereas previous solar neutrino detectors had been sensitive only to the expected ν_e flavor, SNO was simultaneously sensitive to this and the total flux of all active neutrino flavors through the use of a deuterium target; this provided a measurement of the solar neutrino flux independent of the neutrino oscillation hypothesis.

SNO was enormously successful, leading to the conclusion that the solar neutrino problem was due to matter-enhanced flavor oscillations, and providing confirmation of the neutrino oscillation hypothesis favored by Super-Kamiokande’s atmospheric neutrino observations. Art McDonald, director of the SNO experiment, was awarded the 2015 Nobel Prize in Physics for these contributions, along with Takaaki Kajita for Super-Kamiokande.

The strengths of the SNO detector that allowed this important result, namely the unique target medium, excellent shielding due to depth, and exceptionally low backgrounds, also make SNO a useful tool for a competitive measurements of the *hep* solar neutrinos and the ν_e component of the DSNB flux.

5.1 Physical Principles

The key advantage of the SNO design was the choice of target medium: heavy water, i.e. $^2\text{H}_2\text{O}$, commonly referred to as D_2O . The detection of Cherenkov radiation in large water detectors had been demonstrated by nucleon decay and light water solar neutrino experiments, but in terms of neutrino interactions, these detectors were sensitive only to elastic scattering of electrons (ES). A heavy water detector also has detection channels corresponding to the W -mediated charged current (CC) and Z -mediated neutral current (NC) neutrino interactions with deuterium, which have a low threshold of a few MeV. The three interactions of interest are summarized in Table 5.1, with corresponding diagrams in Figure 5.2. The NC and ES interaction are sensitive to all neutrino flavors, while CC is available only for ν_e . Although the total cross section is relatively small, the differential ES cross section is strongly directional, with the direction of the final state electron correlated with that of the incoming neutrino. Meanwhile, the electron energy in the CC interaction is strongly correlated with neutrino energy. Figure 5.1 illustrates the directional and spectral characteristics of the CC and ES interactions. The NC interaction is detectable via the final state neutron, which carries no information about the initial neutrino energy or direction.

SNO was sensitive to DSNB ν_e primarily through the CC interaction. Super-Kamiokande, on the other hand, searches for DSNB $\bar{\nu}_e$ through inverse beta decay,

$$\bar{\nu}_e + p \rightarrow e^+ + n,$$

5.1 Physical Principles

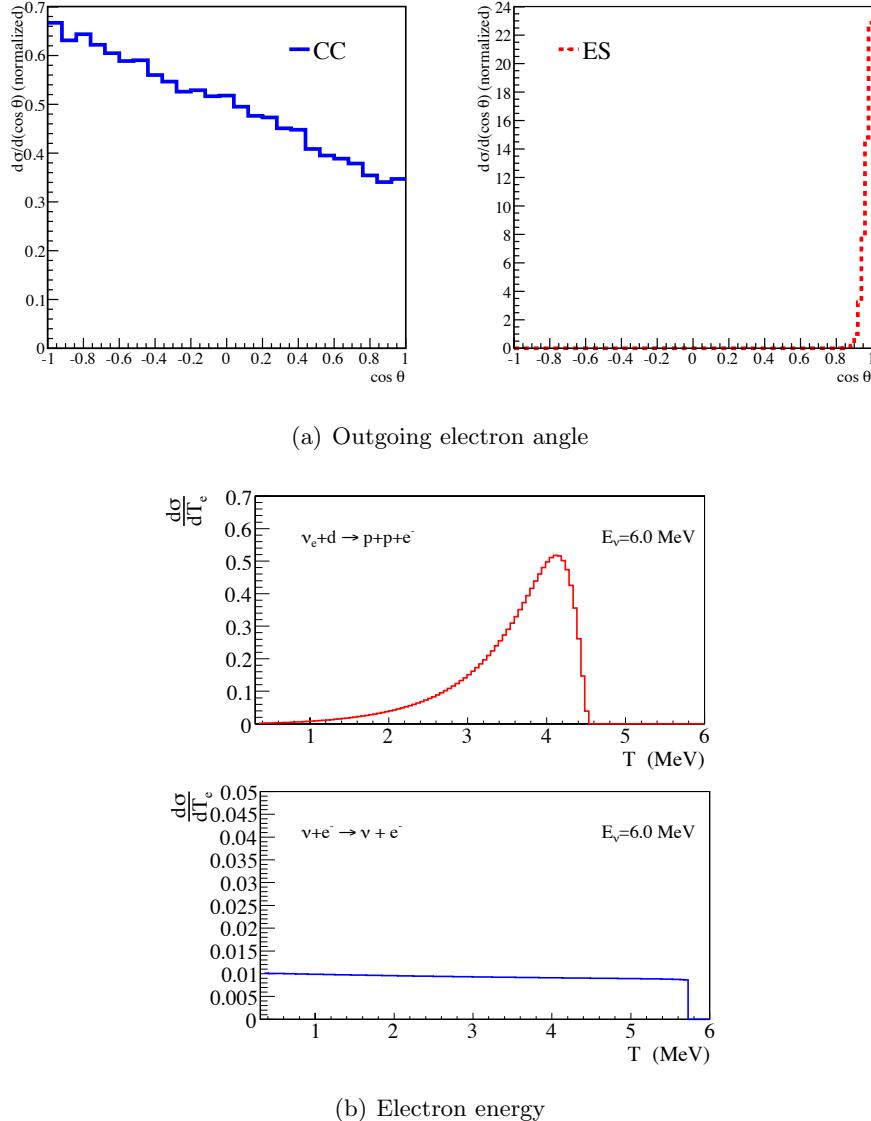


Figure 5.1: Normalized differential cross sections for the $(\nu_e - d)$ CC and $(\nu_e - e^-)$ ES reactions, showing the spectral fidelity of CC and directionality of the ES. Figures reproduced from Reference [27].

5.1 Physical Principles

Elastic Scattering	(ES)	$\nu_x + e^- \rightarrow \nu_x + e^-$	Fig. 5.2(c), 5.2(d)
Charged Current	(CC)	$\nu_e + d \rightarrow e^- + p + p$ ($E_{\text{th}} = 1.44$ MeV)	Fig. 5.2(a)
Neutral Current	(NC)	$\nu_x + d \rightarrow \nu_x + p + n$ ($E_{\text{th}} = 2.2$ MeV)	Fig. 5.2(b)

Table 5.1: Neutrino interactions in SNO.

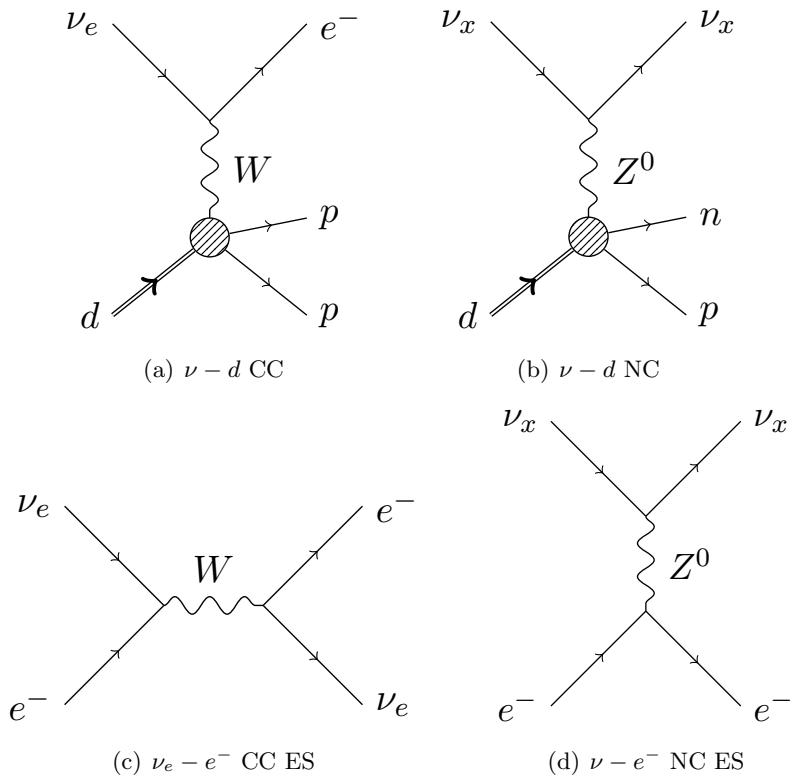
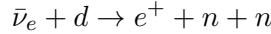


Figure 5.2: Schematic diagrams illustrating the neutrino interactions in SNO.

5.1 Physical Principles

counting positron-neutron coincidences. In addition to IBD, SNO can in principle also observe $\bar{\nu}_e$ through the CC interaction



with background strongly suppressed by requiring a triple coincidence. Unfortunately, given the cross sections and mass of SNO, the event rates expected are very small (see, e.g. Reference [76]); SNO cannot compete with the much larger Super-Kamiokande in detection of the DSNB $\bar{\nu}_e$ signal, hence this search is focused on ν_e which is enhanced by the CC channel.

SNO is able to observe neutrino interactions through the Cherenkov radiation produced by charged final state particles and secondaries. Such detectable particles include electrons produced in CC interactions and scattered in ES, and electrons Compton scattered by de-excitation γ s following the capture of neutrons produced in NC interactions. Cherenkov radiation is produced when relativistic charged particles travel through a medium of index of refraction n with a velocity $v > c/n$ ($\beta > 1/n$); in this case radiation due to the polarization of the material may interfere constructively. The photons form a wavefront with a half angle $\cos \theta = 1/n\beta$ relative to the particle momentum, forming a cone. The imaging of the resulting Cherenkov ring projected onto an array of photon detectors is the fundamental operational principle for ring-imaging Cherenkov (RICH) detectors, providing a means to infer particle track position and momentum. The spectrum of emitted photons per unit wavelength $d\lambda$ and distance dx is given by [28]

$$\frac{dN}{d\lambda dx} = \frac{2\pi\alpha z^2}{\lambda^2} \left(1 - \frac{1}{\beta^2 n^2(\lambda)}\right) \quad (5.1)$$

where α is the fine structure constant and z the charge of the particle in units of the

electron charge. The velocity threshold implies very different energy thresholds for particles of different masses: in water, the total energy threshold for an e^\pm is 0.77 MeV, μ^\pm 160 MeV, π^\pm 210 MeV, and $p \sim 1400$ MeV. Hence, for example, the protons scattered in few-MeV scale CC and NC interactions are not visible in SNO, and below-threshold (“invisible”) μ which decay to isolated relativistic electrons within the detector present a problematic background.

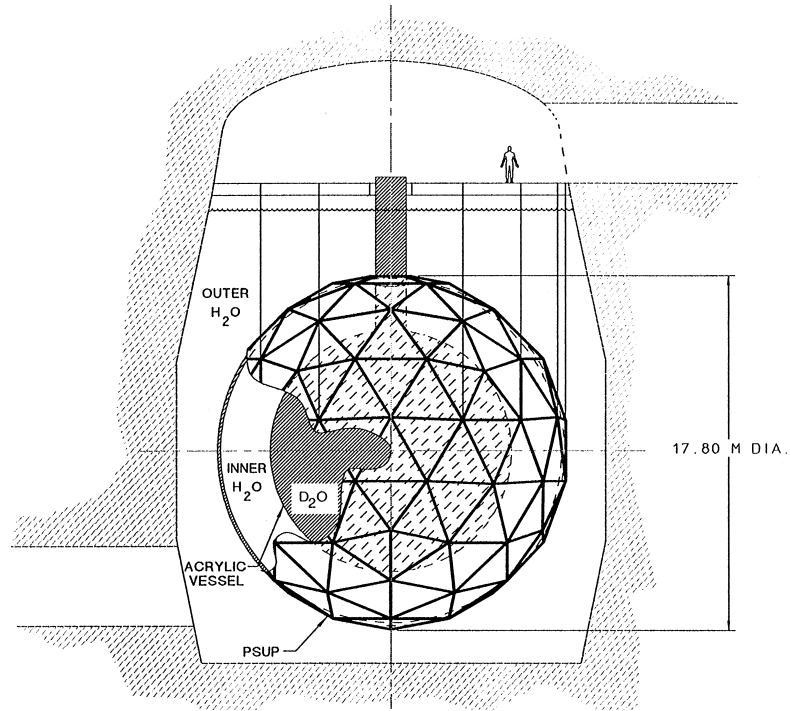
5.2 Detector

The SNO detector is fully described in Reference [77]. The detector consisted of a target volume enclosed within a transparent acrylic sphere 6 m in radius, viewed by 9456 inward-looking 8” Hamamatsu R1408 photomultiplier tubes at a radius of 8.4 m. The acrylic vessel (AV) and the structure supporting the PMTs (PSUP) were suspended in an ultra-pure light water-filled cavity, which was additionally instrumented with outward-looking (OWL) PMTs to provide an active veto system. In order to shield from muons and from the neutrons and decay products of unstable isotopes resulting from muon interactions, the detector was located deep underground with a 6020 m (water equivalent) rock overburden at the 6800-foot level of the Inco⁷ Creighton nickel mine near Sudbury, Ontario, Canada. Figure 5.3 shows the major elements of the detector and the experimental hall.

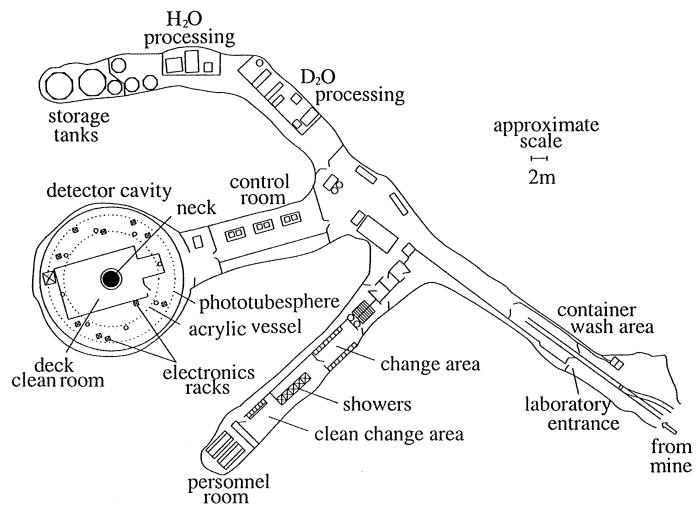
The detector was significantly modified twice throughout the course of the experiment, resulting in three distinct operational phases. These phases differed primarily in the method for detecting neutrons (foremost as the products of NC interactions). In Phase I (the “D₂O phase”), the detector was loaded with only heavy water, and neutrons were detected by capture on deuterium; the de-excitation gammas Compton scattered, producing Cherenkov

⁷This mine is now operated by Vale Canada Limited, and the expanded underground experimental facility by the SNOLAB Institute.

5.2 Detector



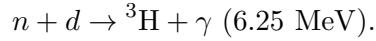
(a) Diagram of the detector.



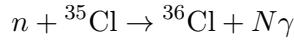
(b) Diagram of the experimental hall.

Figure 5.3: The SNO Experiment. Figures from Reference [77].

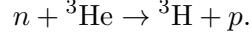
light in the reaction



In Phase II (the “salt phase”), 0.2% NaCl was added to take advantage of the improved neutron capture rate on Cl and the higher-energy and more isotropic de-excitation γ cascade:



with a total energy of the N gammas of 8.6 MeV. In Phase III (the “NCD phase”), an array consisting of 36 ${}^3\text{He}$ proportional counters (referred to as Neutron Counting Devices, or NCDs) plus four additional detectors used for background studies, was deployed in pure D_2O to further improve neutron detection by observing the capture on ${}^3\text{He}$ to produce a triton and proton:



5.3 Electronics

The SNO data acquisition system (DAQ) was responsible for the capture and storage of PMT hit information (threshold-crossing time and integral charge for each PMT fired), including the trigger system which makes a decision on when to record data. The DAQ electronics were comprised of custom-designed printed circuit boards organized into a hierarchical system with many identical boards at each level. PMTs connected to the front-end electronics via a single RG-59 coaxial cable, which carried both the high voltage and any signal pulse; the signal pulse was picked off by a decoupling capacitor at the PMT Interface Card (PMTIC).

Each SNO PMTIC handled 32 PMT channels, arranged into groups of eight which shared a Paddle Card, the physical interface for the cable connections. In addition to HV supply

5.3 Electronics

and signal pickoff, the PMTIC also allowed paddle cards to be disabled through remotely-programmed relays, and included a “charge injection” calibration system for simulating PMT pulses. From the PMTIC, the PMT signal traveled to a corresponding Daughterboard, one of four mezzanine cards on the 32-channel Front-End Card (FEC), each of which handled 8 channels.

The Daughterboards contained discriminators, integrators, and a custom control ASIC for each channel which was responsible for data storage and readout and generation of trigger signals. When a PMT signal crossed a user-defined threshold, the discriminator fired and initiated (a) the start of a timer (TAC), (b) the integration of charge (Q) on capacitors with a long (L) and short (L) integration time and high (H) and low (L) gain, and (c) the generation of a set of current pulses that were fed into the trigger system.

Readout began if the front-end electronics received a Global Trigger (GT) signal from the trigger system. The GT stopped the TAC timer, providing a measurement of a channel’s hit time relative to the detector-wide GT. Integrated charges (QHS, QHL, and QLS or QLL) were stored on a switched capacitor array to await readout. If a GT did not arrive before a fixed channel reset time (~ 400 ns), the data were discarded. Asynchronously, the sequencer (a Xilinx FPGA on the FEC) stepped through channels with data available, piping the analog voltages into a set of analog-to-digital converters (ADCs) shared FEC-wide, and writing the digitized words into FEC memory along with any error condition bits, forming a 96-bit PMT Bundle. The set of these PMT bundles corresponding to a GT constitute the raw channel-level detector data for an event.

The full system consisted of 9728 channels: 19 electronics crates on the deck above the detector each held sixteen 32-channel FECs. 9456 of these channels were used for the inward-looking PMT array, with a smaller number for outward-looking PMTs, test PMTs, PMTs in the neck, and electronics calibration channels. An overview of the system is shown in Figure 5.4.

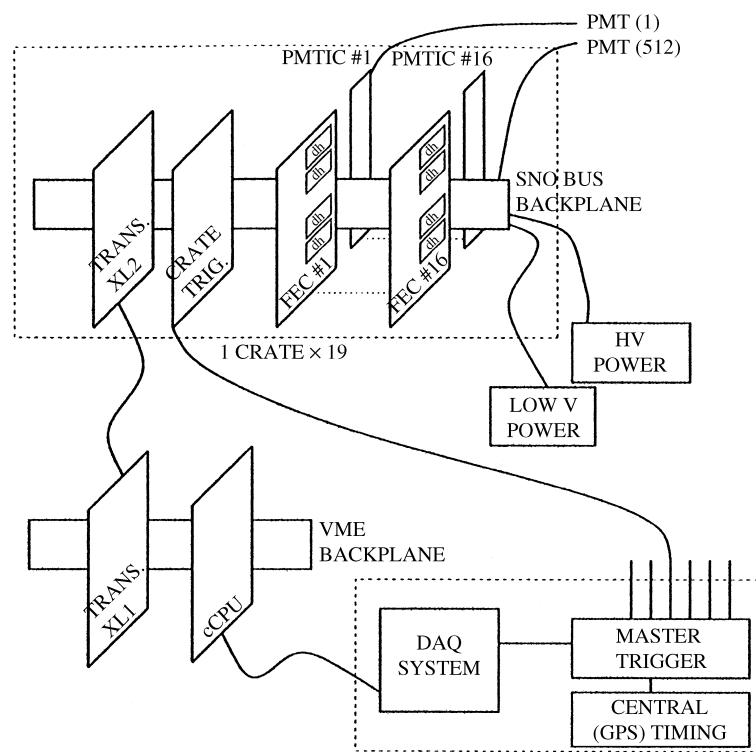


Figure 5.4: A simplified overview of the SNO data acquisition electronics. Figure from Reference [77].

The data acquisition computer was a VME single-board computer, with the FEC memory of the entire detector mapped to local memory via a two-stage translation system: a set of XL1 cards sharing the VME crate were connected to XL2 cards in each front-end electronics crate that translated between VME and the custom SNOBUS protocol (implemented to reduce noise due to readout within the crates). Having been read from the front end electronics, data was passed along to the Event Builder software, which grouped PMT and trigger information according to the sequential Global Trigger ID (GTID) and wrote the sorted data to disk for analysis.

The trigger system was responsible for deciding when to issue a GT and record data, and also had a tiered architecture. When a PMT signal crossed threshold, the individual channel ASIC generated three current pulses: a 93 ns-wide square pulse (NHIT100) and a 20 ns-wide square pulse (NHIT20) of about $600 \mu\text{A}$ each, and a shaped copy of the PMT pulse (ESUM, in a high-gain and low-gain version). These signals were added together with like signals DB-wide, then FEC-wide, then crate-wide on a Crate Trigger Card (CTC). From the CTC, crate-wide trigger sums traveled via RG-58 coaxial cables to the appropriate Analog Master Trigger Card (MTC/A); there were seven MTC/A cards, for NHIT100, NHIT20, low-gain ESUMLO, high-gain ESUMHI, OWL NHIT100, OWL ESUMHI, and OWL ESUMLO. The MTC/A created a detector-wide trigger sum and compared it to up to three user-defined thresholds. Threshold crossings were communicated to the Digital Master Trigger Card (MTC/D) which would issue a GT depending a user-specified mask of enabled trigger channels. The MTC/D also provided a built-in pulser for random sampling of the detector occupancy (pulsed GT), electronics calibration functions, and synchronization with a GPS master clock (connected to a satellite receiver above ground via a fault-detecting, delay-compensated fiber optic system). A schematic overview of the trigger system is provided in Figure 5.5.

The effect of the MTC/A triggers was to provide a threshold setting that scaled with the energy deposited in the detector (the number of photons) within two (100 ns and 20

ns) coincidence windows corresponding to the transit time of light across the detector and from near the center. The pulser was used to create a random trigger at 5 Hz, sampling the detector independent of physics triggers to provide a measurement of e.g. PMT noise rates.

For each GT issued by the MTC/D, trigger conditions were stored and included with the event data. This included the trigger channels causing the GT, the active trigger mask, event times (for both precise inter-event timing and GPS synchronization), ESUM pulse characteristics (peak, integral, and derivative) measured by the Analog Measurement Board (AMB), and error flags. The Event Builder combined this event-by-event trigger metadata with the PMT data to produce complete event records for processing and analysis.

5.4 Calibration

Calibration of the both the detector electronics and the detector response as a whole is essential for understanding the data. Using circuitry built into the electronics and well-understood deployed sources, relative measurements of hit and event energy and time are mapped to absolute measurements in physically-meaningful units.

5.4.1 Electronics

The SNO front-end electronics measured the channel hit times and integrated charges in terms of ADC counts, and small differences in analog circuitry introduced channel-to-channel variations which must be accounted for. The first step in electronics calibration was the ECAL, which consisted of a suite of tests to determine the optimum values for channel-level hardware settings. For example, the slope of the TAC ramps were tuned by setting DACs on the FEC, and the correct DAC values – those which will make the timing measurement consistent across the detector – were measured by the ECAL. The ECAL tests

5.4 Calibration

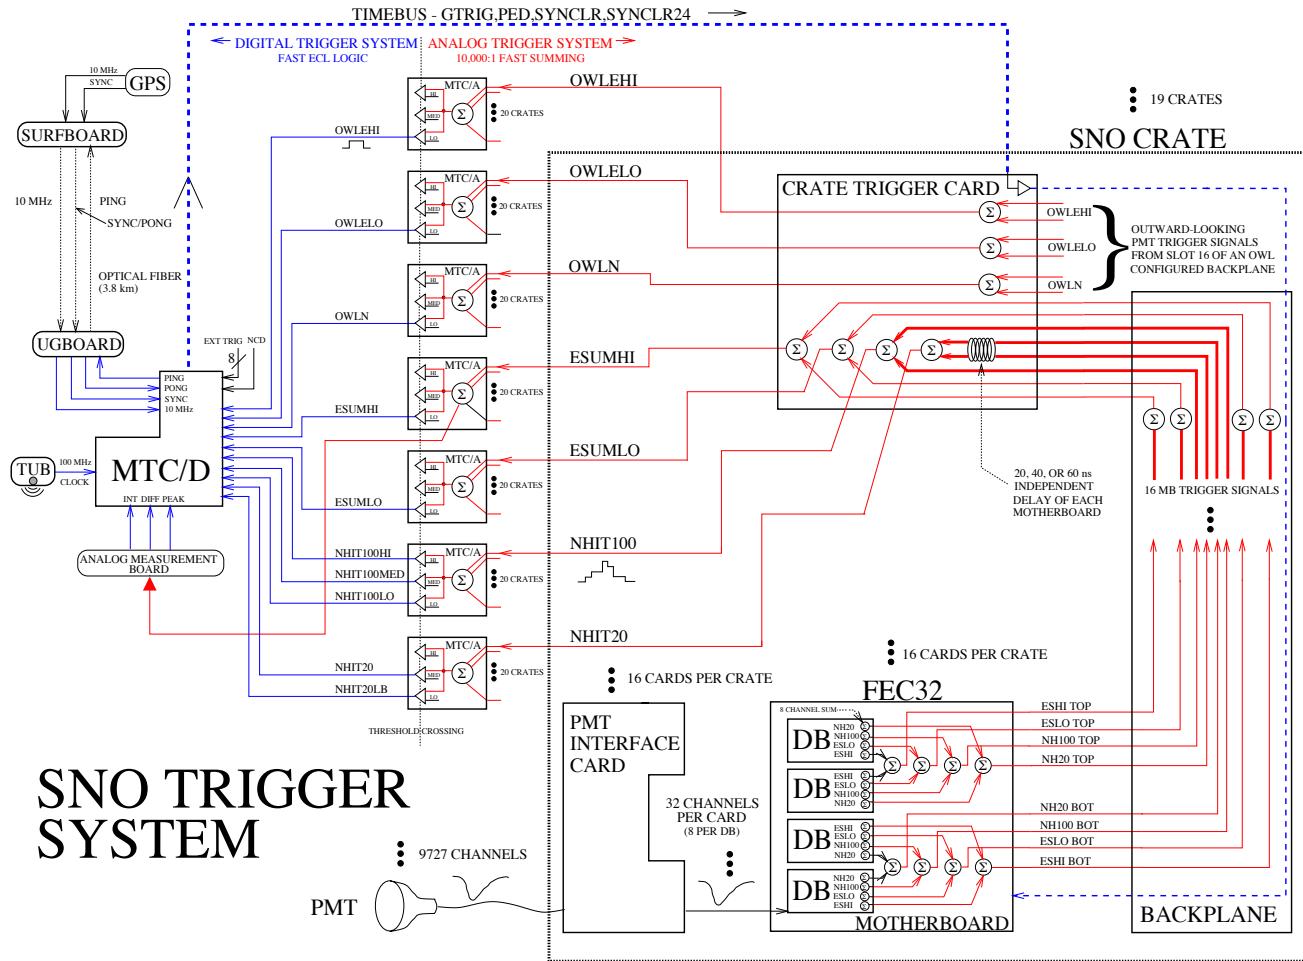


Figure 5.5: An overview of the SNO trigger system. Figure from Reference [78].

5.4 Calibration

made use of the calibration systems built into the trigger system. The so-called “pedestal” (PED) signal forced selected channels to fire, bypassing the discriminator. This provided a measurement of the zero level of the charge integrators, and combining a PED with a precisely-delayed global trigger (GT) allowed calibration of the channel timing. The result of the ECAL was to make the channel-level response as uniform as possible across the detector by tuning low-level hardware settings. ECALs were time-consuming and had to be performed without PMT high voltage, and so were generally executed when electronics were powered up.

The next step of electronics calibration was the ECA, which determined the mapping between the raw ADC counts measured by the front end and physical units required for analysis. The ECA consisted of a PED run which measured the zero-level charges for each channel (which should be relatively uniform due to the ECAL) and a time slope measurement, which used a delayed global trigger precisely calibrated in ns to determine the ADC-to-ns conversion for each TAC. ECA calibrations were performed approximately twice weekly, or following changes to the detector hardware.

Finally, the PCA (PMT calibration) measured the gain, channel-to-channel timing delay (due, e.g. to the path length differences in PCB traces across the FEC backplane), and “time walk” (the dependence of measured time on total charge due to the constant-threshold discriminator bias). The PCA was done using the laserball, a deployed source described in the following section, to generate short pulses of variable intensity. PCA calibrations were performed approximately monthly.

The ECAL constants were loaded to hardware as part of the initialization, and play no further role in the analysis. The constants measured by the ECA and PCA calibrations were applied to the raw data during production processing, to yield time and charge measurements for compensated for electronics effects and suitable for physics analysis.

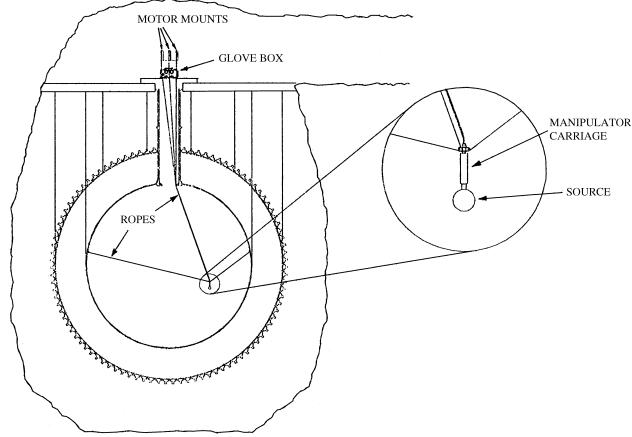


Figure 5.6: The SNO source manipulator system. Figure from Reference [77].

5.4.2 Sources

To understand the response of the detector to the interactions of interest, a variety of radioactive and optical sources were deployed within the target volume. This is essential, for example, to calibrate the absolute energy scale in MeV, rather than the number of PMTs hit, to measure the uncertainty on that energy measurement, to understand the spatial variations across the detector volume, to measure the optics of the media *in situ*, and generally to validate the Monte Carlo detector model.

The deployment of sources was controlled by the manipulator system, a computer-controlled positioning system. While the system could not access arbitrary positions, it was able to sample the $x - z$ and $y - z$ planes (using two positioning ropes) and the z axis (using a single rope) with a position uncertainty of about 2 cm. The concept for off-axis source positioning is shown in Figure 5.6. Beyond the control ropes, sources were generally connected to services in the deck clean room through an umbilical cable; this contained low- and high-voltage cabling, gas supply and return, etc., depending on the requirements of the particular source.

5.4 Calibration

The primary deployed sources included:

Laserball The laserball was a diffusing optical source which isotropically distributed short pulses of light from a multi-wavelength laser system situated on the deck above the detector [79]. This source was used for PMT calibrations (described in the previous section) and in an off-axis configuration to extract the optical properties of the detector media in a dedicated multidimensional fit.

^{16}N The ^{16}N source served as an important energy calibration near the detector threshold [80]. A deuterium-tritium source above the detector produced the gaseous ^{16}N , which was flowed down to a decay chamber in the source. The ^{16}N underwent β decay to an excited state of ^{16}O , which relaxed via the emission of a 6.13 MeV γ . A small PMT scintillator detector inside the source triggered on the initial β , thereby creating a tag for analysis. This signal-like source was also instrumental to understanding the sacrifice of analysis cuts.

^8Li The ^8Li source was, like the ^{16}N , a tagged source using decays in a radioactive gas [81]. The β spectrum is similar to that of ^8B with an endpoint at 14 MeV, and detected through scintillation light produced in the He carrier gas, wavelength-shifted to match the PMT sensitivity by a thin layer of tetraphenyl butadiene (TPB). ^8Li has a short lifetime of 840 ms, requiring the gas — produced by a dT generator and ^{11}B target outside the detector — to be pumped down to the decay chamber at high velocity.

$p\text{T}$ The pT source consisted of a miniature proton accelerator and produced 19.8 MeV γ rays through the $^3\text{H}(p, \gamma)^4\text{He}$ reaction [82]. This high-energy source provided a powerful handle for constraining the energy calibration and a direct measurement of the energy response near the *hep* analysis endpoint. This source was not deployed in Phase II or III, however.

5.5 Simulation

^{252}Cf The ^{252}Cf source was a source of neutron bursts, due to spontaneous fission decays [27]. Gammas were also produced in association with the desired neutrons.

AmBe The Americium-Beryllium (AmBe) source provided neutrons with associated 4.14 MeV γ s, through α decay of ^{241}Am and subsequent $^9\text{Be}(\alpha, n)^{12}\text{C}^*$ [83].

^{238}U and ^{232}Th Acrylic-encapsulated ^{238}U and ^{232}Th provided a higher-rate measurement of U- and Th-chain decays that are important backgrounds to the lower-energy ^8B solar ν analysis [84].

Rn and ^{24}Na Spikes Two distributed sources were used, where a spike of radioisotope-rich water was injected into the detector to provide a uniform distribution of low-energy background without the optical complications of a source container. In separate instances, an Rn spike was used to measure low-energy backgrounds, and a ^{24}Na spike as a distributed 2.6 MeV γ source [27].

The *hep* and DSNB regions of interest are at relatively high energy, with electron kinetic energies $T_{\text{eff}} > 14$ MeV, and so the most relevant sources are ^{16}N , pT , and ^8Li . The laserball, of course, is critical to calibrating the detector in the first place. For the high-energy analysis, tagged Michel electrons provide an additional calibration source, since the spectrum is well-modeled. The use of these calibrations to understand uncertainties in the energy response is detailed in Chapter 7.

5.5 Simulation

This analysis, like previous SNO analyses, relies heavily on Monte Carlo simulations of signals and backgrounds with a detailed detector model to extrapolate calibration data to signal expectations and to characterize systematic uncertainties. All simulations are performed using SNOMAN (SNO Monte Carlo and Analysis), a Fortran 77 software package

developed by the SNO collaboration over the course of the experiment. In addition to simulations, SNOMAN was also used in processing of detector data, for example the application of calibrations and the reconstruction of event vertices. The goal of the SNOMAN simulation is to reproduce detector data as accurately as possible, using microphysical modeling and minimal tuning.

SNOMAN is capable of producing realistic samples of all major signals and backgrounds through a set of flexible event generators. Modeling of solar neutrino generation is included (using the spectrum measured by Winter et al. [85]), as are total and differential CC, NC, and total cross sections for all neutrino flavors and detector materials. Radioactive background generators are provided with tables of β and γ decay chain branching fractions. The DSNB signal is treated similarly to a solar neutrino source, with a spectrum determined as outlined in Chapter 4. SNOMAN also has a capability to bypass event generation and input primary Monte Carlo vertices directly into the Monte Carlo Particle List (MCPL) banks; this is the method used to import atmospheric neutrino interaction vertices from NUANCE into SNOMAN for propagation through the full detector model and reconstruction.

5.5.1 Physics and Detector Modeling

The modeling of light production by electron and γ interactions is handled by the EGS4 (Electron Gamma Shower) package [86]. Cherenkov photon directions are linearly interpolated between discrete track segments used to approximate multiple scattering, with a step size optimized to reproduce data and number of photons is sampled from a Poisson distribution based on the analytic calculation of light yield. Neutron transport is handled by the MCNP4A (Monte Carlo Neutron Propagation) code developed at Los Alamos National Laboratory [87], configured to propagate neutrons only; accurate modeling of thermal neutron propagation through H_2O and D_2O is crucial for understanding the NC signal and the deuterium photodisintegration background at low energy. Higher-energy leptons (such

5.5 Simulation

as cosmic ray muons) are handled with the LEPTO 6.3 [88] package developed at CERN, and hadrons by FLUKA [89] and GCALOR. Lower-energy secondaries produced in these processes are passed to EGS4 for further propagation. This hybrid Monte Carlo model represented the state of the art during SNO running, and is used for the present analysis as well, in order to maintain consistency with past results and make use of existing verification work.

The optical properties of the PMT and concentrator array are complex and handled specially. A partially-phenomenological model is used which employs a full three-dimensional geometry and *ex situ* efficiency measurements scaled in order to reproduce the angular and timing distributions observed in data. If a photon hit is registered in the PMT simulation, a photoelectron is sampled using a model (the electron optics and amplification process are not simulated), and passed to a complete simulation of the detector electronics. The output of this stage is identical to detector data and undergoes the same processing and reconstruction.

5.5.2 Event Rates

For the solar neutrino signals, the expected event rate is the product of the standard solar model flux and survival probability described in Chapter 2, the interaction cross section, and the interaction-dependent detector response which depends on the operational phase. Formally [38],

$$R_{T,i} = N\Phi_i \int_0^\infty \phi(E_\nu) P_{ee,i}(E_\nu) \int_0^\infty \frac{d\sigma(E_\nu, T_e)}{dT_e} \int_T^{T+\Delta T} \frac{dR(T_e, T')}{dT'} dT' dT_e dE_\nu \quad (5.2)$$

5.5 Simulation

for event classes where the energy is measured (e.g. CC, ES), and

$$R_i = N\Phi_i \int_0^\infty \phi(E_\nu) \sigma_{\text{tot}}(E_\nu) dE_\nu \quad (5.3)$$

when the energy is not measured (e.g. NC). Here N refers to the number of targets, Φ_i the total flux for neutrino reaction i with a normalized spectrum ϕ , P_{ee} the ν_e survival probability, σ the cross section for the interaction channel under consideration, T the reconstructed kinetic energy and ΔT the bin width in T , and dR/dT the distribution of outgoing electron energies.

Cross Sections SNOMAN, and hence this analysis, uses $\nu - d$ cross sections for the CC and NC processes calculated by Butler, Chen, and Kong [90], and the $\nu - e^-$ ES cross sections by Bahcall [91]; both calculations include radiative corrections. The total cross sections up to 35 MeV are shown in Figure 5.7.

Survival Probability The survival probability for P_{ee} for the ${}^8\text{B}$ and *hep* solar neutrinos is calculated in the full three-neutrino theory as outlined in Section 3.2, using the Sun-Earth Large Mixing Angle Adiabatic approximation (SELMAA) as developed and used in the SNO combined three-phase analysis [38, 48].

Fluxes The ${}^8\text{B}$ and *hep* solar neutrino fluxes are based on the BS05(OP) standard solar model, as given in Section 3.1. The DSNB flux is calculated following Beacom and Strigari as outlined in Chapter 4. The rate of background events due to atmospheric neutrino interactions is calculated as in Section 6.2.

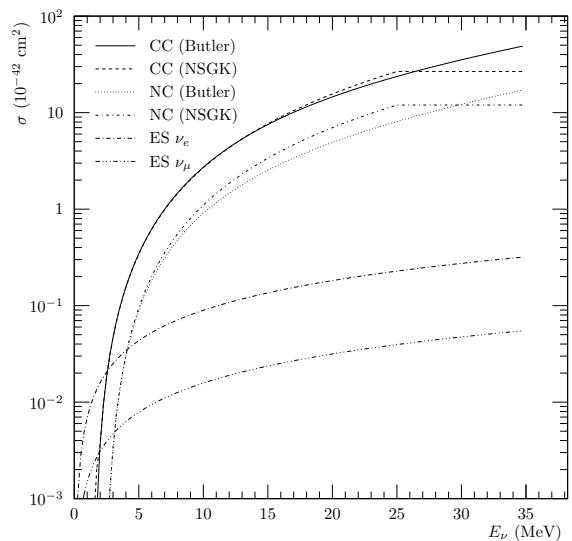


Figure 5.7: Total cross sections for $\nu_e - d$ CC, $\nu - d$ NC, and $\nu - e^-$ ES.

5.6 Vertex Reconstruction

A variety of sophisticated vertex and energy reconstruction algorithms were developed for the SNO experiment, which together take the observed pattern of hit PMTs for an event in charge and time, and calculate the most probable event position, direction, and energy, under the assumption that the event is due to a single electron-like vertex. The vertex fitters chosen for this analysis are the Path Fitter (FTP) [44, 92] for Phases I and II, and the NCD-aware version of the QPDF fitter (nFTU) [47] for Phase III. In all three phases, the Energy Response (RSP) fitter [93] is used to reconstruct the energy, based on the results of the FTP or nFTU vertex fit.

These algorithms are discussed in substantial detail in the aforementioned references; here their general structure is outlined briefly.

5.6.1 The FTP Vertex Fitter

For Phases I and II, the Path Fitter (FTP), fully described in References [44, 92], is used for reconstruction of event vertices (position, direction, and time relative to the trigger). As a first pass, the detector volume is coarsely binned and a grid scan is performed to find the voxel where the likelihood for the time-of-flight corrected PMT hit times (time residuals) as compared to a Monte Carlo-derived PDF is maximized [44]:

$$\log \mathcal{L} = \sum_{i=1}^{N_{hit}} \log [P(t_i - t_e - |\vec{r}_e - \vec{r}_i| n^* / c)] \quad (5.4)$$

where t_i and t_e refer to the positions of the PMT i and the event hypothesis and n^* is an effective index of refraction which determines the average photon group velocity. This simple first-pass fit provides a starting point (seed) for the event position \vec{r}_e and time t_e for the Path Fitter.

The Path Fitter computes the most likely direction \vec{u}_e for the vertex in addition to \vec{r}_e , and uses both the hit times of PMTs and the event-PMT angle, encoding the average angular distribution of Cherenkov photons in a PDF for $\cos \alpha$. The event PDFs are composed of contributions from direct light (D) and other (O) sources of hits, which are subsequently factorized into uncorrelated time (T) and angle (A) PDFs. The total likelihood function is [44]

$$\log \mathcal{L} = \sum_{i=1}^{N_{hit}} \log \left(\sum_{j=D,O} \prod_{k=T,A} P_j^k(\vec{r}_e, \vec{u}_e, t_e; t_i, \vec{r}_i) \right) \quad (5.5)$$

where the time PDF is as above and the angular PDF is a function of the Monte Carlo-derived average Cherenkov angle distribution and the solid angle given the hypothesized event position and direction.

There are a few notable limitations of this approach. The factorization of PDFs assumes that the photons are uncorrelated, which is not true for Cherenkov photons which propagate, but are not produced, independently. Also, the angular PDF is generated by summing many simulated events. In reality, there are significant event-by-event variations in hit angle (due for example to multiple Coulomb scattering) and any particular event is unlikely to look like the average. Finally, there is a systematic ‘drive’ along the event direction introduced as a result of photons scattering out of the Cherenkov cone; a correction is applied to compensate for this, and a systematic uncertainty is included for this ‘vertex accuracy.’

5.6.2 The nFTU Vertex Fitter

In Phase III, the impact of the NCD array on photon trajectories must be accounted for in the vertex fitter. The Path Fitter relies heavily on analytic response functions, and is not easily modified to account for the more complex optics in this configuration. Instead, an NCD-aware version of the QPDF (FTU) fitter, denoted nFTU, is used for vertex reconstruction in Phase III. This fitter is described in Reference [47].

The first step is the maximization of a likelihood function of the same form as Equation 5.4, fitting for the event position \vec{r}_e and time t_e only. In nFTU, the PDFs for the time residual distribution depend on the NCD shadowing as a function of position; they are constructed by interpolating between fully-shadowed and non-shadowed Monte Carlo PDFs, which were additionally truncated to minimize the contributions of NCD reflections. The event direction \vec{u}_e is fit in a later step, with the position fixed, by maximizing a likelihood function depending on the average angular distribution of hits about the hypothesis direction for a Cherenkov ring, with a factor to compensate for solid angle.

5.6.3 The RSP Energy Fitter

While simply counting the number of hit PMTs provides an approximate estimate of the energy of an event, energy reconstruction algorithms incorporate information about the detector and the event vertex (position, direction, and time) to compute (ideally) the true energy of an event, in physical units and independent of the detector. For this analysis, the Energy Response fitter (RSP), described in Reference [93], is used.

RSP operates on PMT hits only in a prompt time window of ± 10 ns in the hit time residuals. An effective N_{hit} is calculated accounting for the number of in-window hits N_w and expected number of accidentally coincident dark (noise) hits N_d , and a position-dependent scaling is applied to map to an effective number of hits for an event at the center of the detector. For an event position \vec{r}_e and direction \vec{u}_e [44],

$$N_{corr} = (N_w - N_d) \times [(\epsilon_r/\epsilon_0) \epsilon_h \epsilon_d]^{-1} \quad (5.6)$$

where ϵ_h and ϵ_d are corrections for the number of channels online (with good hardware and calibration status) and a time-dependent variation in the response (“drift”), and ϵ_r characterizes the optical response, with ϵ_0 for the center. The factor ϵ_r is a weighted

sum of angular response functions over event directions (the event position is fixed) and wavelengths. This accounts for the wavelength distribution of Cherenkov light and the average attenuation in the optical media (D_2O , acrylic, and H_2O), and also includes a correction factor for multiple photons striking the same PMT.

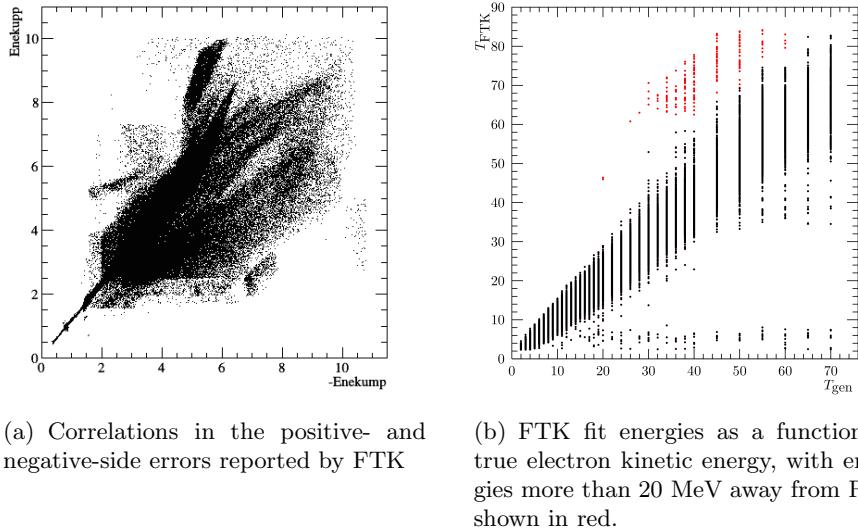
Subsequent to the Phase I SNO publications, substantial improvements were made to the RSP algorithm by MacLellan [93]. In the improved version, used in this analysis, the angular response function in ϵ_r uses the response characteristics of individual PMTs, rather than averaging over a large angular bin. This serves to improve the reproduction of local variations in the response and better capture the effects of multiple incident photons.

With the variation in time, angle, and position accounted for, the final step is to put the output in terms of physical energy units. To achieve this, N_{corr} is computed and averaged for monoenergetic electron Monte Carlo, and this is repeated for many electron energies. The resulting mapping is inverted to obtain an energy in MeV from the RSP estimate. This provides the electron-equivalent reconstructed kinetic energy, denoted T_{eff} , used throughout this work.

5.6.4 Issues with the FTK Energy Fitter

FTK is a newer, more sophisticated energy fitter used heavily in the SNO low-energy threshold analysis [46]. The energy resolution for FTK appears to be slightly better than RSP for events in the *hep* and DSNB ROI, making this a promising potential improvement; particularly, any reduction in the 8B background is a boon to the *hep* search. FTK also reports asymmetric uncertainties which can be used to cut events that fit poorly and are likely to be misreconstructed.

Unfortunately, while FTK is very successful for events with $T_{\text{eff}} \lesssim 20$ MeV, at higher energies there is significant misreconstruction which cannot be cut by the uncertainties, as these errors become highly irregular and do not scale linearly with resolution. It is thus



(a) Correlations in the positive- and negative-side errors reported by FTK

(b) FTK fit energies as a function of true electron kinetic energy, with energies more than 20 MeV away from RSP shown in red.

Figure 5.8: Examples of FTK mis-reconstruction of high-energy events, shown for Phase I isotropic electron Monte Carlo after corrections.

not possible to construct an FTK uncertainty cut with reasonable signal loss (less than a few percent) for high-energy events. These issues are illustrated in Figure 5.8. Not only are these the events of interest for the DSNB search, but high-energy Michel electron events are used in validating the reconstruction in the *hep* and DSNB energy range, and so we are forced to abandon FTK for this high-energy analysis. In principle, there is nothing wrong with the algorithm, and it remains possible that further development could yield a robust high-energy fitter with better resolution than RSP.

These issues are clearly evident when looking at Michel electrons, used in this analysis as a high-energy calibration source, above the DSNB search window. Figure 5.9 compares the performance of fitters in reconstructing stopped μ Monte Carlo; RSP outperforms FTK significantly.

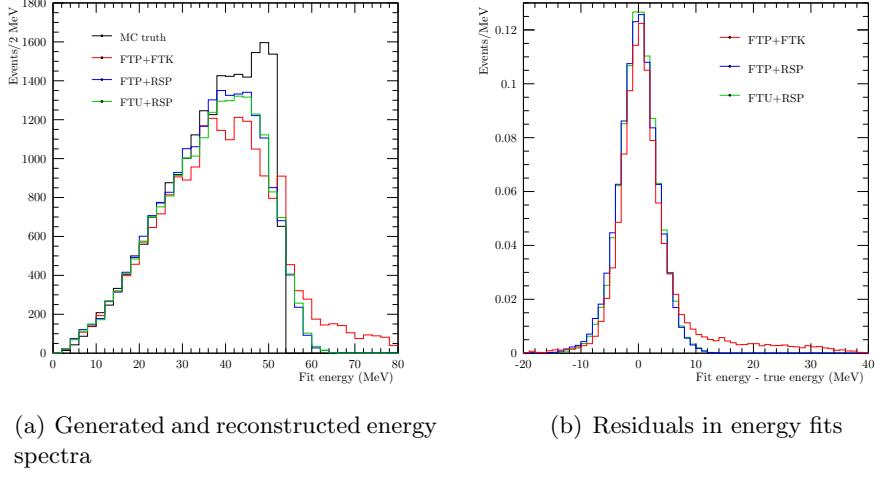


Figure 5.9: A comparison of SNO fitter performance for reconstructing Phase I stopped μ Monte Carlo.

5.7 SNO+ Upgrades

This section describes detector upgrades made for the SNO+ experiment, a successor to SNO primarily concerned with a sensitive search for neutrinoless double-beta decay ($0\nu\beta\beta$). The physics aspects of SNO+ are described in Appendix A.

The SNO+ detector, shown in Figure 5.10, makes use of the existing infrastructure of the SNO experiment described above, but with the target volume filled with Tellurium-loaded liquid scintillator rather than heavy water. To account for the buoyancy of scintillator-filled inner vessel in the water-filled cavity, a hold-down rope net system has been constructed [94]. Additionally, to meet more stringent background requirements, an embedded LED and laser calibration system has been installed, reducing the dependence on a deployed optical source for calibration of the PMT timing and *in situ* measurement of the scintillator optical properties [95].

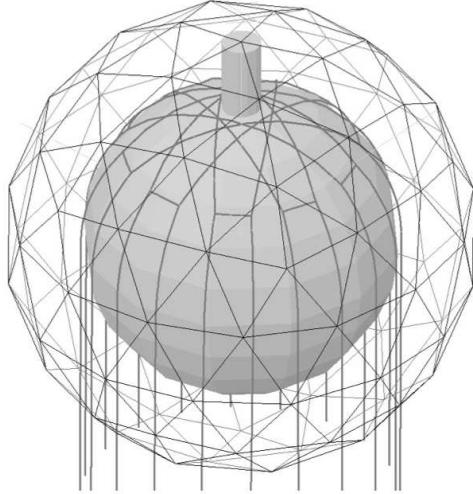


Figure 5.10: The SNO+ detector. The active volume is contained within a 12 m diameter acrylic sphere (grey), suspended with ropes in a volume of ultra-pure water. The outer 18 m diameter structure supports 9500 inward-looking PMTs.

5.7.1 Electronics Upgrades

The transition from a water Cherenkov to a liquid scintillator detector significantly changes the requirements for the front-end and trigger electronics described in Section 5.3. Both the rate (the number of events per second) and the occupancy (the number of channels hit per event) are expected to be much higher for SNO+, necessitating faster front-end readout and a replacement of the analog trigger system which has limited ability to handle the higher RMS current resulting from a “brighter” detector. The hardware trigger threshold for SNO+ has not yet been determined, but the intention is to take data as quickly as the hardware allows, and potentially perform data reduction in a software (“Level 2”) trigger. This approach facilitates low-energy time-correlated analysis, such as looking for proton scatters in a supernova burst, or low-energy α decays to tag background β s.

In order to improve the data transfer rate, the performance of electronics calibrations, and the level of operator control, we have upgraded the front-end readout with a new crate

controller board, the XL3, which functionally replaces both the XL1 and XL2. An XL3 is responsible for the control of a single crate and autonomously and continuously reads out data from front-end cards. Data is packaged into ethernet packets and pushed to a central DAQ computer, then passed onto an Event Builder similar to the system built for SNO. The logic on the XL3 is implemented in a Xilinx Virtex-4, a device which combines a PowerPC CPU and an FPGA, coupled with a high-speed shared data bus. The ethernet communication is handled in C code using the lwIP (lightweight IP) library running on the PPC, and the SNOBUS protocol communication is implemented in the FPGA. This system improves on the SNO readout speed by about a factor of 100, by leveraging faster hardware and by parallelizing the work of crate readout. The total bandwidth is about 300 Mbps, sufficient for rates expected with SNO+ calibration. The proof-of-concept and initial versions of the C-layer data buffering and IP communication were implemented by the author.

I have also developed, installed, and commissioned an upgraded the analog master trigger card (MTC/A) with a drop-in replacement board dubbed the MTC/A+. The primary goals were to eliminate the limitations in trigger sum current inherent to the MTC/A and to expand the dynamic range beyond the ~ 1000 -hit range of SNO, allowing us to trigger at thresholds anywhere in the full 10000-hit range, and also to capture trigger sum waveforms with the maximum dynamic range.

The required analog changes presented an opportunity to make several additional upgrades. While the MTC/A+ is entirely compatible with the existing SNO trigger system, it was fully redesigned to improve threshold stability, provide additional diagnostic capabilities, and incorporate programmable logic to allow hardware triggering on certain signals of interest. These improvements will have a direct impact on the live time and physics capabilities of the SNO+ experiment.

Analog Design The analog side of the MTC/A+ was designed to optimize speed and dynamic range. In particular, a dynamic range covering the full spectrum of 1–10000 hit channels is required, as well as better than single-hit resolution at low hit levels, fast response (pulse rise and fall times of less than 3 ns per hit channel), high-rate performance (better than 10 kHz at 100 hit channels), and long-term (> 1 hour) drifts in the DC baseline of less than 1 PMT hit. The MTC/A+ addresses these criteria with a two-stage operational amplifier sum in each of three gain paths and through active baseline restoration.

The gain paths are tuned to saturate at approximately 1000, 5000, and 10000 hit PMTs, with the “high-gain” (low NHIT threshold) path offering the best resolution in NHIT and the “low-gain” (high NHIT threshold) the possibility of triggering on and digitizing the full 10000-channel scale. A block-level diagram of the MTC/A+ is shown in Figure 5.11.

Good signal integrity — low rise, fall, and settling times and low noise — is maintained by using high-speed current-feedback amplifiers in an inverting configuration. The particular amplifier used in the MTC/A+, the Texas Instruments THS3001, was chosen for its fast slew rate and wide power supply range (± 15 V). The large signaling current of the CTC presents a challenge for downstream electronics. Each hit channel results in $600 \mu\text{A}$ of current in the NHIT100 CTC sum, meaning that a fully populated crate sources 300 mA, and a 10000-hit event results in 6 A arriving at the analog trigger. The MTC/A featured a current limit at the input, which dumped excess current and limited its range to about 2000 PMT hits. Given the higher number of hit PMTs expected in SNO+, this is not a viable option for the MTC/A+. Instead, low-noise amplifiers are used to sample the voltage of an attenuated input signal, while still maintaining an adequate signal-to-noise ratio.

The MTC/A+ also performs active baseline restoration. In the fully DC-coupled SNO system, long-term drifts in the input current effectively raised or lowered the trigger threshold, necessitating frequent recalibration. An active *LRC* feedback loop restores the MTC/A+ baseline to the nominal value with a time constant on the order of tens of seconds. This not only eliminates the impact of slow thermal fluctuations in the crate current output, but

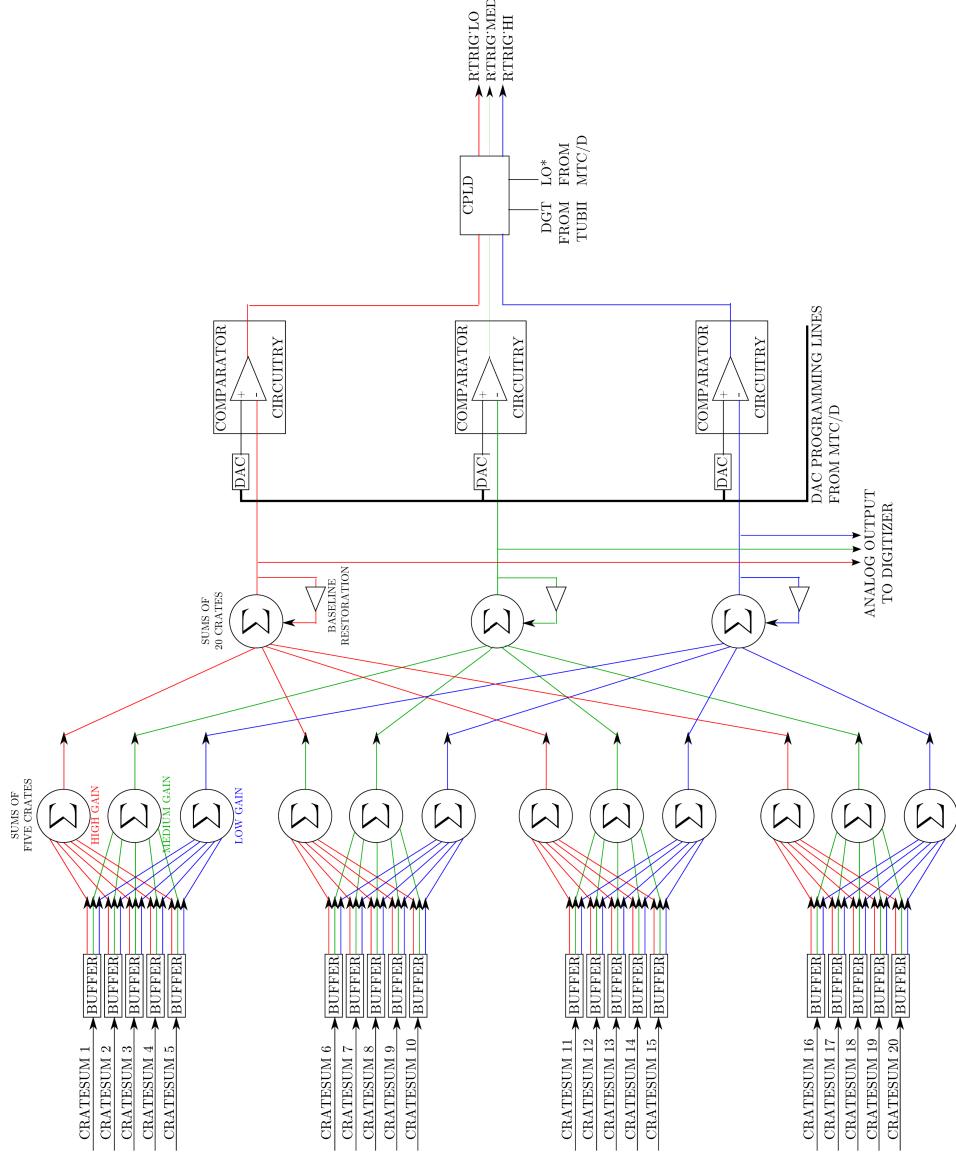


Figure 5.11: A block-level overview of the major features of the MTC/A+, the upgraded analog trigger board for SNO+.

also helps to address the issue of trigger “dropout.” Dropout is an issue with the trigger current pulse generation in the front-end CMOS chips, resulting in a channel trigger that stays high until the next time the channel is hit, effectively lowering the trigger threshold by one hit in the interim. The MTC/A+ baseline restoration time constant attempts to strike a balance where we compensate for dropped-out channels, but do not integrate real signals.

Trigger Logic The trigger logic on the MTC/A+ is implemented in VHDL code running on a Xilinx XC2C512 CPLD. The most basic form defines the following logic:

1. Fire raw trigger on channel X when channel X crosses threshold
2. Also fire raw trigger on channel X at the end of the lockout window⁸ if channel X crosses threshold during lockout window, or is still high at the end of the lockout window
3. Fire raw triggers at the end of N sequential lockout windows, where N is a programmed number of “forced retriggers.”

The CPLD may be reprogrammed if additional logic is desired. For example, it is possible to use the three separate copies of the trigger signal to create a gated low-energy window trigger, where an initial high-energy trigger enables triggering within a low-energy window for some fixed amount time. This could be used to create a hardware trigger for Bi-Po coincidences, improving tagging of an important background for the SNO+ $0\nu\beta\beta$ search. “Bi-Po” refers to the decay chains $^{214}\text{Bi} \rightarrow ^{214}\text{Po} + \beta \rightarrow ^{210}\text{Pb} + \alpha$ ($t_{1/2} = 164.3 \mu\text{s}$) and $^{212}\text{Bi} \rightarrow ^{212}\text{Po} + \beta \rightarrow ^{208}\text{Pb} + \alpha$ (299 ns). The noted half-lives for the α decay are such that the β and α may fall in the same 400 ns trigger window (“pileup”) or not. The MTC/A+ can be configured such that an initial β enables a gated α trigger, even if the hardware threshold

⁸The lockout window is a ~ 440 ns period after the MTC/D issues a global trigger during which new any triggers are ignored; the PMT data for this window is essentially already flagged for saving.

is set much higher in energy. This ensures we capture these coincidences, reducing the risk of missing an α and subsequently misidentifying a lone β as a $0\nu\beta\beta$ signal event.

Operational Improvements In addition to the analog and trigger logic changes, new MTC/A+ features also improve trigger system operations. Each crate-level input passes through a remotely-programmable reed relay, so that an operator can disconnect a crate from the sum without physically moving any trigger system cabling. This provides a useful debugging tool, and can improve detector live time in the case where restricted underground access prevents fixing or physically disconnecting a problematic crate.

Status The XL3 and MTC/A+ boards have been installed in the SNO+ detector since 2010 and used in all subsequent detector electronics and DAQ commissioning runs, and have performed well. In particular, the production MTC/A+ boards have a resolution of about 3 DAC counts per hit PMT on the high-gain channel, low noise, and very stable thresholds. Complete testing of trigger efficiency under realistic conditions, however, awaits deployment of the laserball source during the initial water-filled phase of SNO+.

Chapter 6

An Improved *hep* and DSNB Search with SNO

Given the unique sensitivity of the SNO experiment to the *hep* solar neutrinos and ν_e component of the DSNB, it is vitally important that searches for these signals are performed using the full data set comprising all three phases, expanding on the existing Phase I-only analysis [15] and previous studies of fits for the *hep* flux [16]. This chapter provides an outline for such a three-phase analysis and highlights improvements I have made relative to previous work.

6.1 Signals and Backgrounds

As the regions of interest for the *hep* and DSNB searches, where the backgrounds are lowest, occurs at relatively high energy, the low-energy backgrounds that were of central importance to previous analyses (such as the low-energy threshold analysis [46]) are of much lesser significance. The dominant backgrounds are due to interactions of ^8B solar neutrinos and atmospheric neutrinos. There is also potential for other backgrounds if the fiducial volume is expanded beyond the nominal 550 cm, such as events attributed to light production in the AV acrylic [45].

The signals and backgrounds considered for this analysis include:

- *hep* ν_e CC on deuterium
- *hep* ν_e CC on ^{17}O
- *hep* ν_e CC on ^{18}O
- *hep* ν ES in the target region
- *hep* ν NC on deuterium
- DSNB ν_e CC on deuterium
- DSNB ν ES in the target region
- ^8B ν_e CC on deuterium
- ^8B ν_e CC on ^{17}O
- ^8B ν_e CC on ^{18}O
- ^8B ν NC on deuterium

6.2 Atmospheric Neutrino Backgrounds

- ${}^8\text{B}$ ν ES in the target region
- ${}^8\text{B}$ ν ES in the AV acrylic
- Atmospheric ν interactions
 - All ν flavors, $E_\nu > 100$ MeV
 - ν_e , $E_\nu < 100$ MeV
 - $\bar{\nu}_e$, $E_\nu < 100$ MeV

The energy spectra for these backgrounds (after application of corrections and cuts described below) are shown in Figure 6.1, for Phase I. Other sources of neutron captures are not distinguished from solar neutrino NC interactions; the total NC PDFs are used to represent all neutrons.

6.2 Atmospheric Neutrino Backgrounds

Primary cosmic rays (primarily protons) interact in Earth’s atmosphere to produce secondary mesons that can eventually decay to neutrinos, which are known as atmospheric neutrinos. The energies of these neutrinos span several orders of magnitude, resulting in a broad array of observable processes in detectors. Atmospheric neutrinos are a rich subject in their own right, and provided the first strong evidence for neutrino oscillations when an asymmetry in the upward- and downward-going fluxes was found to be consistent with an oscillation model by Super-Kamiokande in 1998 [19]. For the purposes of this analysis, however, atmospheric neutrinos are considered only as a background, with a rate determined by the absolute flux. These interactions result in a significant background for the *hep* search which becomes the dominant background in the DSNB region.

There are large uncertainties in the absolute fluxes — on the order of 25% — due to uncertainties in measurements of the primary cosmic ray spectra and the hadronic interactions

6.2 Atmospheric Neutrino Backgrounds

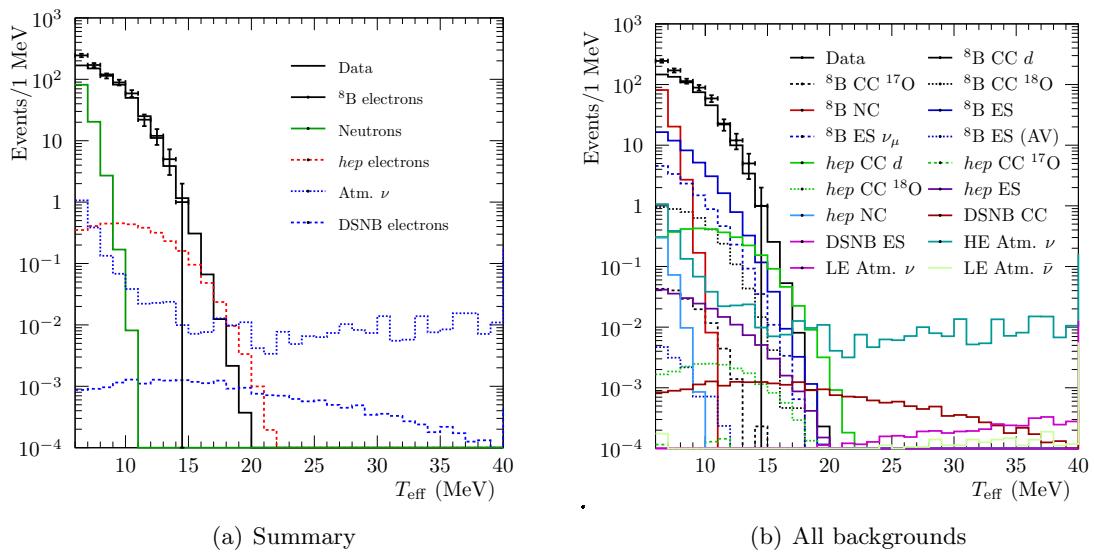


Figure 6.1: Energy spectra around the *hep* and *DSNB* energy regions of interest for Phase I, after corrections and cuts described in the text.

6.2 Atmospheric Neutrino Backgrounds

in the atmosphere. Unlike oscillation searches that use ratios to largely cancel systematic uncertainties, we are here forced to propagate this large error to the expected background level in the *hep* and DSNB search windows. Additional systematics related to interactions in the detector also play a role, as described below.

6.2.1 Production

For neutrino energies below ~ 100 GeV, the primary interactions generating neutrinos in the atmosphere are due to secondary pions:

$$\pi^\pm \rightarrow \mu^\pm + \overset{\leftarrow}{\nu_\mu} \quad (6.1)$$

$$\mu^\pm \rightarrow e^\pm + \overset{\leftarrow}{\nu_e} + \overset{\leftarrow}{\nu_\mu}. \quad (6.2)$$

At higher energies, kaons come to dominate the flux.

In terms of the slant height for a height h given a density ρ , $X(h) = \int_h^\infty \rho(h')dh'$, the flux ϕ_j of a cosmic ray of type j is given by the set of coupled cascade equations of the form [21]

$$\frac{d\phi_j(E, X)}{dX} = -\frac{\phi_j(E, X)}{\lambda_j(E)} - \frac{\phi_j(E, X)}{d_j(E, X)} + \sum_k S_{k \rightarrow j}(E, X) \quad (6.3)$$

where λ_j is the interaction length which depends on the cross section, d_j is the decay length calculated from the lifetime, and $S_{k \rightarrow j}$ gives the flux of j due to a parent k . The first term describes disappearance due to attenuation, the second particle decay, and the third is a source term. This equation is analytically tractable only through substantial approximation. However, some interesting features are evident, for example that for unstable secondaries there exists a critical energy where $\lambda_j \sim d_j$, i.e. where interaction effects dominate decay. For pions, $E_{\text{crit}}^{(\pi)} = h_0 m_\pi / \tau_\pi \sim 115$ GeV; above this threshold, kaon decay dominates neutrino

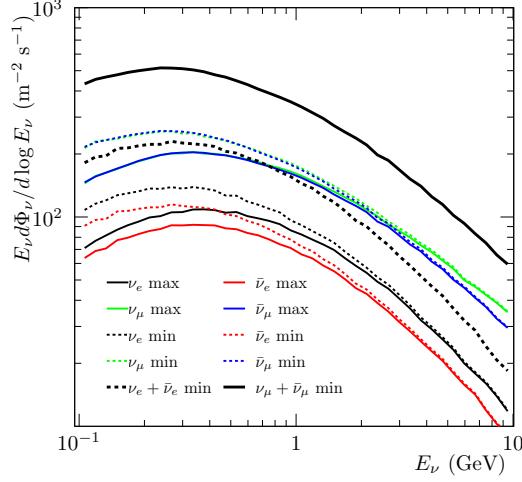


Figure 6.2: Fluxes calculated in the Bartol04 calculation for each flavor, averaged over zenith angle. Min and max refer to the solar minimum and maximum. The summed spectra for $\nu_e + \bar{\nu}_e$ and $\nu_\mu + \bar{\nu}_\mu$ at the solar minimum are shown for comparison to Figure 2 in [96].

production⁹.

More detailed calculations are carried out numerically, and modern approaches include a full three-dimensional treatment with charged cosmic rays being deflected by the Earth's magnetic field, which affects the angular distributions particularly at low energies. For the purposes of this analysis, I have used the Bartol04 calculation of the atmospheric neutrino flux [96] from 100 MeV to 10 GeV. In that work, the authors calculated differential fluxes for Kamioka, Soudan, and SNO; the fluxes relevant for SNO are shown in Figure 6.2. At lower energies ($E \leq 10$ GeV), the solar wind affects the cosmic ray flux, introducing a modulation with the solar cycles. The SNO data set spans from 2001 to 2006, from the maximum of solar cycle 23 to the minimum between cycles 23 and 24. The Bartol04 fluxes are calculated for the solar minimum and maximum; I use the average in calculating the

⁹Here, $h_0 \sim 6.4$ km, the decay constant of atmospheric density assuming an exponential model.

6.2 Atmospheric Neutrino Backgrounds

	Flux ($\text{m}^{-2} \text{s}^{-1}$)	
	Solar Minimum	Solar Maximum
ν_e	14824.0	10973.1
$\bar{\nu}_e$	12157.4	9452.8
ν_μ	27853.1	21304.5
$\bar{\nu}_\mu$	27967.4	21283.6

Table 6.1: Total fluxes of atmospheric neutrinos predicted in the Bartol04 model, integrated over energy and zenith angle.

total flux. The total flux used to normalize the expected high-energy event rate is calculated from the Bartol04 predictions by integrating over zenith angle and energy for each flavor. These rates are given in Table 6.1.

For the lower-energy fluxes below 100 MeV, tables computed by Battistoni et al. [97] were used.¹⁰ In this regime, only ν_e and $\bar{\nu}_e$ are considered, as these events are below the muon production threshold and the ES contribution is very small. Fluxes are calculated only for the solar minimum, where the rate is highest; this leads to an overestimate of this background, which is neglected since the sub-100 MeV contribution to the total atmospheric neutrino background is small.

6.2.2 Oscillations

Neutrinos produced in the atmosphere undergo oscillations which, as described in Section 2.2.1, results in a survival probability that depends on the energy and the path length; the

¹⁰The authors calculated fluxes for Gran Sasso and Kamioka only; the Sudbury tables were provided to the SNO collaboration by request and do not appear in the publication.

6.2 Atmospheric Neutrino Backgrounds

observed flux will be related to the initial flux as

$$\phi_\beta'(E_\nu, \theta_z) = \sum_{\alpha=e,\mu} \phi_\alpha(E_\nu, \theta_z) P_{\alpha\beta}^{3\nu}(E_\nu, L(h(E_\nu, \theta_z), \theta_z)), \quad (6.4)$$

which is computed numerically. Here, ϕ_α represents the initial flux for a neutrino flavor α as a function of energy E_ν and zenith angle θ_z , $P_{\alpha\beta}^{3\nu}$ is a transition probability for flavor α to β in a three-neutrino mixing model, and L is the propagation path length, which is a function of the production height h . The distribution of path lengths has been calculated by Gaisser and Stanev [98] as a function of neutrino energy and zenith angle using a 3D model. I take the Bartol04 ν_e and ν_μ flux predictions, and for each energy/angle bin, sample a large set of neutrino energies and zenith angles θ_z . Production heights h are sampled from the appropriate Gaussian distributions, defined according to the Gaisser and Stanev calculation, and oscillation baselines calculated using the relation [21]

$$L = \sqrt{(R_E + h)^2 - (R_E - d)^2 \sin \theta_z} + (R_E - d) \cos \theta_z \quad (6.5)$$

with R_E the radius of Earth and d the detector depth. This method averages over neutrino energy, zenith angle, and production height within the bin. Finally, oscillations are applied according to a three-neutrino oscillation model with the best-fit oscillation parameters given in Section 2.3. The mean survival/appearance probability for the ensemble computed for each bin is used to reweight the Bartol04 neutrino (antineutrino) flux in that bin, leading to a suppression of the ν_μ ($\bar{\nu}_\mu$) flux and the appearance of a ν_τ ($\bar{\nu}_\tau$) flux, with a small impact on $\langle \bar{\nu}_e \rangle$. The oscillated fluxes are given in Table 6.2.

In order to normalize the Monte Carlo simulations, I have calculated event rates for each neutrino flavor in each detector volume using the oscillated Bartol04 fluxes and the CC, NC, and total cross sections for detector materials extracted from GENIE [99]; the

6.2 Atmospheric Neutrino Backgrounds

	Flux ($\text{m}^{-2} \text{s}^{-1}$)	
	Solar Minimum	Solar Maximum
ν_e	14382.53	10650.25
$\bar{\nu}_e$	11796.31	9174.87
ν_μ	15119.12	11670.56
$\bar{\nu}_\mu$	15166.70	11631.49
ν_τ	13640.63	10266.80
$\bar{\nu}_\tau$	13629.04	10274.74

Table 6.2: Total oscillated fluxes of atmospheric neutrinos predicted by the Bartol04 model and the oscillation model described in the text, integrated over energy and zenith angle.

		ν_e	$\bar{\nu}_e$	ν_μ	$\bar{\nu}_\mu$	ν_τ	$\bar{\nu}_\tau$
Total	Solar Min.	87.27	23.66	92.45	31.55	18.91	8.62
	Solar Max.	73.51	20.65	81.96	28.05	16.07	7.32
CC	Solar Min.	64.90	15.53	66.55	19.95	0.35	0.18
	Solar Max.	54.69	13.61	59.25	17.95	0.35	0.18
NC	Solar Min.	22.36	8.12	25.95	11.59	18.56	8.43
	Solar Max.	18.81	7.03	22.74	10.10	15.72	7.13

Table 6.3: Expected number of atmospheric neutrino events per year in the heavy water during Phase I.

implementation of the SNO detector in GENIE is discussed in Appendix B. The expected event rates for the D₂O volume in Phase I are given in Table 6.3. On average, a total of 262.5 (685.8) events are expected within the D₂O (full detector) per year assuming the flux at solar minimum.

6.2 Atmospheric Neutrino Backgrounds

6.2.3 Atmospheric ν Backgrounds to the *hep* and DSNB Search

Given the high energies of atmospheric neutrinos, many channels are available for interactions in detectors: quasielastic (QE) scattering, production of light mesons (e.g. π), light baryons (e.g. $\Delta + \pi$), strange baryons (e.g. $\Sigma + K$) in inelastic scatters on nucleons, and elastic and quasielastic scattering on electrons. Fortunately, many of these interactions involve multiple decays and coincidence tagging may be used to substantially reduce the background. The dominant backgrounds for the *hep* and DSNB searches in SNO are 15.1 MeV γ rays that result from NCQE interactions on ^{16}O and from Michel electrons resulting from the decay of sub-threshold μ produced in CCQE (and to a lesser extent resonant pion production) interactions.

In the γ case, an initial $\nu(^{16}\text{O}, ^{15}\text{O}^*)n$ or $\nu(^{16}\text{O}, ^{15}\text{F}^*)p$ reaction leaves the struck nucleus in an excited state. A few percent of these states ultimately relax via the emission of a 15.1 MeV γ from $^{12}\text{C}^*$ [100], although the branching ratio is not well measured. A large systematic uncertainty is included to account for this.

The Michel electrons due to “invisible” muon decay form the dominant background above the *hep* endpoint (the DSNB region of interest). The Cherenkov threshold for a μ in water is 160 MeV (see Section 5.1), and so water Cherenkov detectors are blind to low-energy μ produced in CCQE and interactions where resonantly-produced sub-threshold pions decay to sub-threshold μ . If these unseen μ decay within the detector, the resulting relativistic electrons will appear as isolated single-electron events, mimicking the signal.

Other interactions are generally reducible, either via coincidence tagging or properties of the events themselves, such as the shape or number of Cherenkov rings. These approaches are described in Section 6.3.2.2 for the present analysis. For a description of the similar cuts used in the Super-Kamiokande DSNB search, see Reference [14].

For the low-energy ($E_\nu < 100$ MeV) atmospheric neutrinos, only the dominant ν_e and $\bar{\nu}_e$ CC interactions are simulated. The former looks essentially like a high-energy, isotropic solar neutrino signal, while the latter produces coincident neutrons via inverse beta decay.

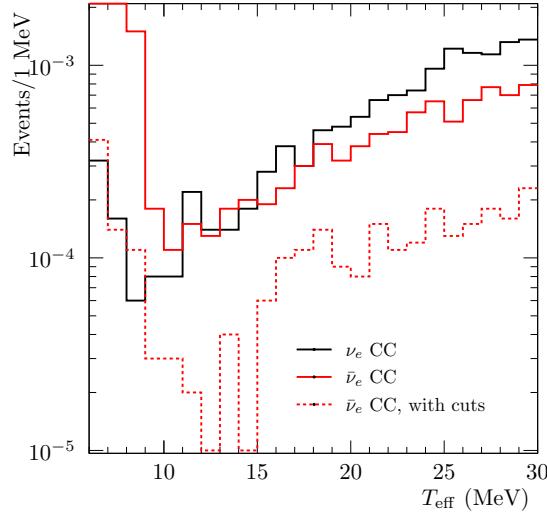


Figure 6.3: Energy spectra for low-energy ($E_\nu < 100$ MeV) atmospheric neutrino CC interactions in Phase I.

The $\bar{\nu}_e$ component is therefore significantly reduced by coincidence cuts, as shown in Figure 6.3.

6.2.4 Systematic Uncertainties

The estimation of flux and cross section uncertainties follows that of Reference [15]. The flux uncertainty, based on comparisons of primary cosmic ray flux measurements and different theoretical approaches, is approximately 10% for $E_\nu < 10$ GeV [101]. This depends on energy, and we conservatively use a constant 10%. Cross section uncertainties are defined separately for CCQE interactions (25%) and all other interaction types (30%) based on Reference [102]. An uncertainty of 100% is assigned to the production rate for 15.1 MeV γ rays following NCQE interactions on oxygen nuclei.

Finally, a detailed comparison of atmospheric neutrino Monte Carlo to data in sideband control regions found that neutron production was under-predicted in the model [15, 103]. Therefore the MC is reweighted based on the presence of a coincident neutron in order to better match data, and a systematic uncertainty is applied to these weighting factors. The factors are 2.11 ± 0.32 for events with neutrons and 0.55 ± 0.04 for those without neutrons [103].

6.2.5 Simulations

The primary vertices (positions, momenta, and types of final-state particles) for the high-energy ($E_\nu > 100$ MeV) atmospheric neutrino Monte Carlo were produced using the NUANCE package [104], and then passed into the full SNOMAN detector simulation for propagation through a realistic detector model [103]. For the NUANCE step, a simple model of the SNO detector was implemented — a 5 cm thick spherical shell of $C_5O_2H_8$ acrylic with heavy water inside and light water outside out to 8.5 m — and a Bartol04 fluxes and a nominal neutrino mixing model used as input.

For the present analysis, the latest mixing parameter measurements are used in a three-neutrino mixing model as described above; the NUANCE/SNOMAN MC is reweighted to transform from the assumed oscillation model to our own. For calculation of event rates used in the reweighting, I have used a model based on the state-of-the-art GENIE package [99]. The GENIE model and the implementation of the SNO detector are described in Appendix B.

6.3 Data Selection

There are three stages involved in the selection of events from the dataset to be used in the analysis. First, entire runs (a contiguous block of time, usually several hours, where

the detector configuration is constant) are flagged as good or bad. For the selected runs, individual events are then subject to low-level cuts meant to remove instrumental and time-correlated backgrounds, and finally to high-level cuts which test compatibility with the signal hypothesis.

6.3.1 Run Selection

In order to ensure that only high-quality data is selected for analysis, a set of criteria are applied to each run. For each phase, a set of detector conditions were specified that constituted stable running, and the run selection criteria were designed to flag anomalies. These criteria included data quality metrics automatically calculated by the First Pass (FPS) and Second Pass (SPS) selection tools during data processing, and also information recorded during data taking by the shift operator. The log files from the event builder and the environmental monitoring system (CMA) also bear on the run quality. All of this information was considered by members of the SNO run selection committee on a run-by-run basis. More details on the SNO run selection process, including descriptions of the specific criteria, may be found in Appendix B of Reference [105].

This process culminated with run lists for each phase containing thoroughly vetted high-quality physics data, which were further reviewed before publications. The present analysis makes use of the following run lists:

Phase I (D₂O) The same run list is used as for the SNO NC Physical Review Letter, Reference [43]. This includes runs with abnormally high radon concentration, which were excluded from the subsequent low-energy threshold analysis (LETA) [46]; these runs correspond to about 10% of the Phase I live time. The total live time for Phase I for this analysis is 306.4 days.

Phase II (Salt) The run list used for the SNO Phase II Physical Review C publication [45]

excluded runs with high levels of radon or activation of the sodium¹¹, which present backgrounds for lower-energy oscillation analysis. These runs are added back for the *hep*/DSNB analysis. The salt *hep* run list was created by the authors of the 2006 *hep* analysis in support of Reference [16], and includes 370 more runs than the salt PRC run list, for a total of 1582 runs and a 22.3% increase in live time. The total live time for Phase II is 478.63 days.

Phase III (NCD) The same run list is used as for the SNO combined three-phase Physical Review C publication [48]. The total live time for Phase III is 385.17 days.

For calibration source data, the same run lists are used as are recommended in existing source analysis documentation for the pT source (enumerated in Appendix E), and a run list for ${}^8\text{Li}$ source data was constructed based on past analysis efforts and detector shift reports (see Appendix F).

6.3.2 Event Selection

In order to maximize the sensitivity of the *hep* and DSNB searches, the background in the search region must be minimized. Fortunately, most backgrounds in this high-energy regime are different from signal in fundamental ways: many instrumental backgrounds (due e.g. to electronics effects) have distinctive signatures, and atmospheric neutrino interactions often produce secondaries visible in time-correlated events, or multiple/non electron-like prompt Cherenkov rings. The signal in both searches is single, electron-like Cherenkov rings produced in isolation. Since the energy is high relative to other SNO searches, there is less multiple scattering and improved photon statistics, both of which improve reconstruction and allow tighter cuts.

¹¹Specifically activation of ${}^{24}\text{Na}$, which produces decay products too low in energy to be of concern for the present analysis.

In the following sections, I define a set of event selection criteria, variables constructed from the event observables that discriminate between event classes. For the counting analysis, these parameters are subject to optimized cuts, and the signal region of interest becomes a volume in the hyperspace of observables. For the signal extraction fit analysis, we can make use of the shapes of the distributions of these variables by treating them (including correlations) as dimensions in a probability distribution.

In the following sections, the term “sacrifice” refers to acceptance loss, i.e. the fraction of the signal that is cut by an effort to remove background.

6.3.2.1 Low-level Cuts

Low-level cuts are those that generally do not rely on event vertex reconstruction, and are applied to remove events from the data set prior to analysis, never being treated as observables in a fit. They include instrumental background cuts (which typically rely on the geometry, charge, and timing of hit PMTs) and cuts that are based on the conditions under which the event occurred (such as occurring just after a tagged μ interaction).

Instrumental backgrounds are caused by detector effects, for example high-voltage breakdown of a PMT or electronic pickup. Such events tend to have distinct signatures, such as correlations in the physical locations of electronics channels, which are very different from “physics” events. Over the course of the SNO experiment, a variety of cuts were developed that identify instrumental backgrounds with very good accuracy.

The low-level cuts were applied during processing in SNOMAN, and packed into a set of two 32-bit words known as the DAMN banks, such that analyses may share a set of cuts using a bitmask. For each Phase, I adopt the same set of low-level cuts used in previous work [48], as these have been extensively validated and tuned for minimal signal sacrifice (for the ${}^8\text{B}$ solar neutrino region of interest). A summary of the low-level cuts is given in Appendix D.

6.3.2.2 High-level Cuts

The next step in event selection is to define selection criteria which discriminate the *hep* and DSNB signals from physics backgrounds. For the counting analysis, these are treated as cuts, while for a multivariate analysis or fit, some are used as observables. These cuts, which were developed for previous SNO analyses, have been adapted and retuned for this work.

It is assumed that ${}^8\text{B}$ and *hep* events of the same visible energy are indistinguishable¹²; that is, the difference in visible energy spectrum provides the only discriminant. In addition to a having different spectrum, DSNB events also have no correlation with the direction to the Sun. The cuts described in this section are geared toward reducing the background due to atmospheric neutrino interactions where possible, i.e. other than single-electron final states.

Event Isotropy Signal events — *hep* and DSNB ν_e CC interactions at 15-30 MeV — produce single, electron-like Cherenkov rings, which are highly anisotropic. Two variables were developed in SNO for quantifying event anisotropy, β_{14} and θ_{ij} .

The parameter θ_{ij} is the mean angle between all pairs of PMTs hit in a given event. The charge-weighted extension, Q_{ij} , is defined by

$$Q_{ij} = \frac{\sum \sum_{\text{PMTs } i \neq j} q_i q_j \theta_{ij}}{\sum_{\text{PMTs}} q_i}, \quad (6.6)$$

that is, a normalized sum of angles weighted by the charge observed in each PMT. In the LETA analysis [46] this was found to provide good discrimination for low-energy external

¹²There may in fact exist some possibility for discrimination based on Cherenkov ring properties which we are not leveraging. For example, differences due to multiple scattering and δ ray production, since the average true energy of events beyond the ${}^8\text{B}$ endpoint are different, and these factors are presumably not correlated with the statistical fluctuations Cherenkov in photon production.

Bi and Tl events, and used a cut. However, in the high-energy *hep* and DSNB ROI, the cut is not so powerful: for example, with $\sim 99\%$ signal efficiency, only a small amount (about 3–5%) of the (atmospheric neutrino) background is removed. Therefore Q_{ij} is not used, and β_{14} provides the measure of event isotropy.

The β_{14} isotropy parameter is defined as:

$$\beta_{14} = \beta_1 + 4\beta_4 \quad (6.7)$$

where

$$\beta_l = \langle P_l(\cos \theta_{ij}) \rangle_{i \neq j}, \quad (6.8)$$

an average of the l -th Legendre polynomials evaluated at the cosines of the angles between all pairs of hit PMTs i and j . The particular combination of β_{14} was chosen for its good discrimination power and approximately normal distribution [105].

In-time Ratio (ITR) The in-time ratio (ITR) is the fraction of the hits in the event falling within a narrow (~ 8 ns) prompt timing window. For single-electron signal events, the Cherenkov light arrives at the PMTs within a narrow time window, whereas for events such as atmospheric neutrino interactions, rings from multiple secondaries may pile up in the same event window, leading to a more uniform hit timing distribution. The vertex reconstruction is performed under the assumption of a single Cherenkov ring; whether due to pile-up of multiple events (e.g. secondary pions) in the event window or multiple promptly-generated Cherenkov rings, a departure from this hypothesis results in a poor position fit, leading to broadening of the time-of-flight corrected PMT hit time distribution.

Angular Figures of Merit Two additional cuts depend on the reconstructed vertex position and direction, using Kolgomorov-Smirnov (KS)-like tests to evaluate hypothesis that

6.3 Data Selection

the hits were due to a single electron ring, which is useful for identifying both muon rings (too sharp) and multi-ring events. The first (denoted P_ϕ) tests compatibility of the azimuthal distribution of hits around the reconstructed direction to the distribution expected for an electron. This essentially captures the Cherenkov angle, smeared by multiple scattering as averaged over many events. The second test (P_{2d}) is a two-dimensional extension that also includes the polar angle. This calculation makes use of the same empirical 1D angular distribution, but also accounts for scattering, reflections, and noise with an approximate analytical function.¹³

In the previous Phase I *hep* analysis, these figures of merit were calculated within (and only available for) the FTP fitter and used a single PDF generated for a 5 MeV electron event. Subsequently, prior to the LETA analysis [46], they were generalized to apply to all fitters, and also upgraded to include the energy dependence of the angular PDFs. There remains some residual energy bias, which is corrected by scaling P_{2d} by T_{eff}^4 . These updated versions are used for the present analysis, an improvement over previous work.

These figures of merit form powerful discriminants, particularly when the only events of interest are single-ring electron events. In principle, one could make even further use of these variables in an energy endpoint search, since ${}^8\text{B}$ CC events in the *hep* ROI have leaked into that energy range due to upward statistical fluctuations in the number of PMTs hit. The true electron energy for these events is lower than the average true electron energy for *hep* events in the ROI. On average, therefore, ${}^8\text{B}$ events are subject to more multiple Coulomb scattering, and have greater angular spread within the Cherenkov ring. Unfortunately, even with the energy-dependent PDFs, the effect is much too small to observe in P_ϕ or P_{2d} . However, it remains possible that including this information in the upstream fitter, rather than averaging over the multiple scattering effects, could yet provide some discrimination.

¹³These tests use a binned approximation to the true Kolmogorov-Smirnov test (which relies on sorted, unbinned data) which is valid in the limiting case where the number of events is large compared to the number of bins, and the binning is small relative to any physically-meaningful scale for the parameter.

In-cone Time Another figure of merit computed based on the fit vertex is the in-cone time (ICT) parameter. This is a KS test that compares the time residuals for hits inside the Cherenkov ring (defined by $0.6 < \cos \alpha < 0.8$, where α is the angle between the reconstructed vertex and hit) to a template distribution extracted from ^{16}N calibration source data.

Clean Event Burst Cut The low-level cuts described in Section 6.3.2.1 do not tell the full story of whether an event has occurred in isolation with respect to other physics events, which is a requirement for *hep* and DSNB signal events; specifically, we seek to eliminate any events that occur in coincidence with a neutron capture or Michel electron in order to reduce the background due to inelastic scatters of atmospheric neutrinos. For example, a CC DIS in which a neutron (that is later captured) is produced in coincidence with an invisible μ^- (that subsequently decays) would not be flagged by the standard low-level cut criteria. Based on the cut developed for the Phase I *hep* analysis [15], I define a “clean event burst” cut which is triggered by more than one physics-like event within a few hundred ms window. To trigger the cut, an event must meet the following criteria:

- $T_{\text{eff}} > 4 \text{ MeV}$
- $-0.12 < \beta_{14} < 0.95$
- $\text{ITR} > 0.55$
- Radius $r < 600 \text{ cm}$
- $P_{2d} \cdot E^4 > 10^{-5}$
- $P_\phi > 10^{-9}$
- Passes all low-level cuts except retrigger and muon follower

Once such an event is identified, the cut steps backward 125 ms and forward 125 ms, and if any other “clean” events are found in that window, both those and the trigger event are flagged as a clean event burst and excluded from the analysis. The definition of a clean event is extended to include events where the outward-looking PMTs trigger, and NCD-triggered events in Phase III.

Fiducial Volume The fiducial volume is set to 550 cm for this analysis, which has been the choice for most previous SNO analysis efforts. We maintain this status quo for two main reasons, beyond precedent and ease of comparison: background due to so-called IAVB (isotropic acrylic vessel background) events and energy calibration uncertainties. IAVB events are fast isotropic bursts of light; the origin is not known, but is believed to be due to triboluminescence in the AV acrylic. A more complete description is given in Reference [45]. Previous work on *hep* searches set an upper limit of 0.002 in the range $6 < T_{\text{eff}} < 35$ MeV, indicating that this background is negligible [15]. However, the number of accepted IAVB events increases substantially nearer the AV. Given that these backgrounds are not fully understood, it does not seem prudent to use a Monte Carlo model to predict the rate of this background in the region near the AV in a low-statistics search; it would be difficult to trust any purported *hep* or DSNB discovery based on events in this volume.

A second reason for maintaining the historical 550 cm fiducial volume relates to the calibration facilities in SNO. The source manipulator system was not capable of reaching beyond 550 cm in the $x-y$ plane, and so the only way to sample the volume from 550–600 cm was on the z axis, at the bottom of the detector (or by interpolating between source positions inside and outside the AV). This makes it difficult to evaluate the spatial nonuniformity of position and reconstruction algorithms in the region near the AV wall, and hence to define plausible systematic errors without relying significantly on detector Monte Carlo.

6.4 Blindness

In order to minimize bias in the final results, a form of statistical blindness [106] is applied wherein only a fraction of the data is used to develop the analysis, and all parameters are fixed before analyzing the full data set; no further adjustments are made based on the results. In the present analysis, cuts are tuned exclusively based on Monte Carlo samples and validated using sideband and calibration data, and the visible fraction of the data is used only as a cross-check to ensure that there are no serious problems with the analysis before total unblinding. While strictly speaking this is a gray area (a choice to do something different could have been made based on some of the data), the “check” data is still included in the final analysis dataset, consistent with previous SNO analyses that have used a similar blindness scheme.

Since the entire SNO dataset has already been analyzed, including for a *hep* and DSNB search in Phase I, a truly blind analysis is not possible. Instead an iterative pseudo-blind approach is used. 2/3 of the data are re-blinded, and the remaining 1/3 are initially visible to check the analysis. The non-blinded (“unblind”) events are randomly-sampled contiguous sequences of events representing 1/3 of the livetime of each run. In this way, it is possible to develop cuts that identify time-correlated effects.

Each SNO run (both the Monte Carlo and processed data) was split into 30 “dataset” chunks of equal time, and each event tagged with a dataset word. So as to avoid bias, the starting point for dataset 1 was selected randomly based on the run number. In order to select one third of the live time, 10 of the 30 datasets were chosen at random to arrive at the dataset mask, and the corresponding events were copied into new files, to eliminate any possibility of accidental unblinding. This procedure applied to detector data only; the full set of Monte Carlo was available throughout the analysis.

6.5 Monte Carlo Simulations

This analysis uses SNOMAN Monte Carlo described in Section 5.5 and predominantly uses the production MC generated for the LETA [46] and combined three-phase [48] analyses. Simulations are performed for each background, for each run, with the event rate scaled up by a large factor so as to oversample distributions used to build PDFs.

The version of SNOMAN used for this analysis includes many improvements over that used for the Phase I *hep*/DSNB search, including a substantially improved model of PMT charge and timing, and PMT-specific efficiencies; a summary of the upgrades may be found in Reference [83]. These improvements generally serve to improve the agreement of Monte Carlo with calibration data, and thereby reduce systematic uncertainties.

I have generated Monte Carlo samples for low-energy atmospheric neutrino interactions, i.e. those with $E_\nu < 100$ MeV, as these were not created as part of the standard SNO production running. These simulations were generated using SNOMAN version 5.0294.

6.6 Summary of Systematic Uncertainties

A wide range of systematic effects impact this analysis; for example, variations in the energy response may shift ${}^8\text{B}$ background events into and out of the *hep* energy region of interest in a box analysis, affecting the background estimate, or errors in cross sections may change the interpretation of event rates in terms of fluxes. These may generally be classified in two categories: normalization uncertainties, which affect the overall rate of events of a particular class, and shape uncertainties, which change the how the events are distributed in one or more observable dimensions. Table 6.4 provides an overview of the parameters considered in this analysis.

6.6 Summary of Systematic Uncertainties

Parameter	Variation	Correlated
Live time	§6.3.1	
Energy scale	§7.3	
Energy resolution	§7.3	
Vertex accuracy	2.9%	
Vertex resolution	2.5 cm	
Instrumental cut sacrifice	§8.1.1	✓
${}^8\text{B} \nu_e$ spectrum	§6.6	✓
ν Mixing parameters	§2.3	✓
Atm. ν flux	§6.2.4	
$E_\nu > 100$ MeV	10%	✓
$E_\nu < 100$ MeV	25%	✓
Cross sections		
CC $\nu - d$	1.2%	✓
Atm. ν CCQE	25% (§6.2.4)	✓
Atm. ν other	30% (§6.2.4)	✓
15.1 MeV γ	100% (§6.2.4)	✓
Atm. n multiplicity	7% (§6.2.4)	✓

Table 6.4: Systematic uncertainties. “Correlated” means that the parameter is assumed to take on the same value across all phases.

6.6 Summary of Systematic Uncertainties

Live Time A small uncertainty in the live time (the period of time the detector was fully online and collecting data) is derived by comparing live times calculated by using the global 10 MHz system clock and by using randomly-triggered (pulsed GT) data, and also considering the uncertainty due to electronics effects and burst cuts; see e.g. Reference [45]. The fractional live time uncertainties used in this analysis are taken from previous SNO publications using the same or similar data sets, per Section 6.3.1.

Energy Systematics The reconstruction-related systematics are understood through the use of calibration sources, specifically the ^{16}N source (5 MeV) at the low end, the pT source (19.8 MeV, Phase I only) near the *hep* endpoint, and a sample of Michel electrons extracted from data on the high-energy end; the ^8B -like ^8Li source data is used a cross-check and a high-statistics sample with which to search for unexpected tails in the energy response.

The estimation of energy systematics is detailed in Chapter 7. For Phases I and II, the FTP vertex fitter is used, while FTN is chosen for Phase III. The RSP energy estimator is used for all three phases. For more details on vertex and energy reconstruction, see Section 5.6.

Vertex Reconstruction Systematics As with the energy estimation, uncertainties in the reconstructed position and direction can distort observed spectra and shift events into or out of the analysis window. These uncertainties have been estimated based on comparison of calibration data and Monte Carlo for previous SNO publications; I have used values from Reference [48].

Instrumental Cut Sacrifice The low-level instrumental background cuts described in Section 6.3.2.1 unfortunately will also remove a small number of signal events. This cannot be estimated using Monte Carlo, since instrumental backgrounds are not modeled in the simulation. Instead, an estimate of the sacrifice is based on calibration source data.

6.6 Summary of Systematic Uncertainties

While most past efforts have focused on the energy range relevant to ${}^8\text{B}$ oscillation analysis, in the preparation of Reference [15] an estimation was made for higher energies (up to 35 MeV) using a combination of ${}^{16}\text{N}$, ${}^8\text{Li}$, pT , and laser source data. The instrumental cut sacrifice was fit with a quadratic function, as shown in Figure 6.4. This result is used for the present analysis, and assumed to apply to all three phases.

It is clear that the fit is poor at the high-energy side, specifically the high-energy laser data. The systematic error shown in the figure is about 3% up 8.5 MeV and increasing thereafter to a maximum 11% at 40 MeV. For the present analysis, the errors beyond 8.4 MeV are increased such that the maximum at 40 MeV is 25%, more representative of the spread in calibration data. A simple linear scaling is used to expand the uncertainty σ , where

$$\sigma(T > 8.4 \text{ MeV}) \rightarrow (1.0 + 0.0394 \cdot (T - 8.4))\sigma.$$

${}^8\text{B} \nu_e$ Spectrum Shape The best measurement of the energy spectrum for ${}^8\text{B} \beta^+$ decay neutrinos [85] is inferred from a positron energy spectrum, which is in turn inferred from a measurement of ${}^8\text{Be}$ decay α energies. The difficulty of the measurement and the corrections in each step result in an effective energy scale uncertainty in $P(E_\nu)$. This energy shift is modeled as a distortion in the shape of the spectrum, and Monte Carlo events are reweighted according to the parent neutrino energy. The distortion is shown in Figure 6.5.

Solar Neutrino Oscillation Parameters The oscillation parameters relevant for solar neutrinos, Δm_{12}^2 , θ_{12} , and θ_{13} , have been tightly constrained by recent global analyses combining the results of solar neutrino observatories, KamLAND, and short-baseline experiments. For this analysis the parameter values recommended by the Particle Data Group [28] are used; these are also summarized in Section 2.3. Values are sampled within their respective uncertainties in order to propagate this systematic.

6.6 Summary of Systematic Uncertainties

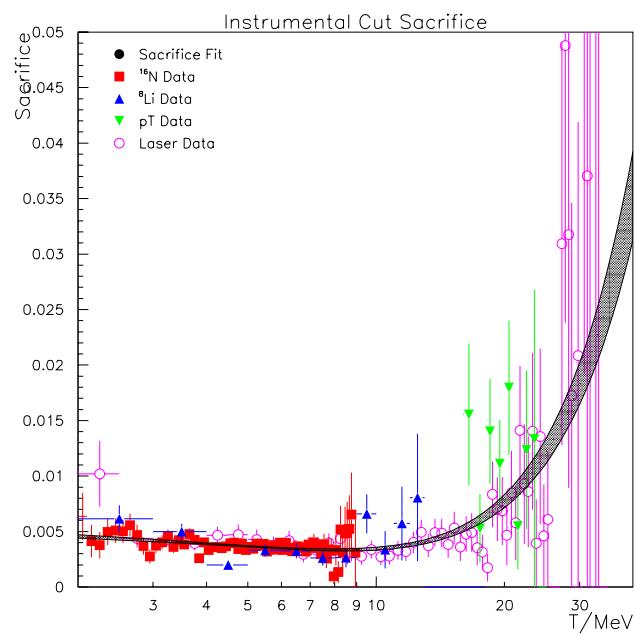


Figure 6.4: Signal sacrifice due to instrumental background cuts, as measured in the Phase I [103]

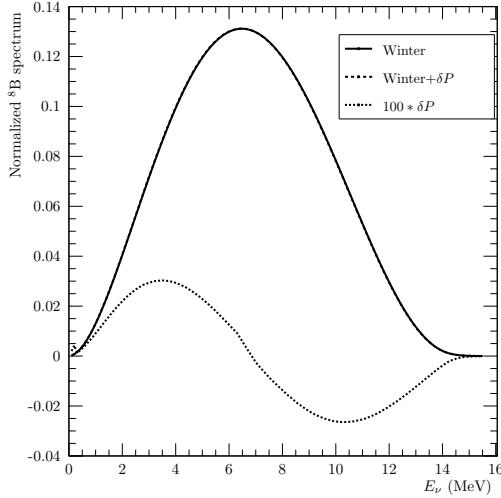


Figure 6.5: The shape uncertainty in the Winter solar ${}^8\text{B}$ ν_e spectrum.

6.7 Corrections to Data and Monte Carlo

The SNOMAN simulation is tuned to match data as accurately as possible, but in some cases final corrections to observables are applied to compensate for detector effects known to be imperfectly modeled, or to improve the agreement of simulations and data.

There are a few corrections applied to the position (FTP) and energy (RSP) observables used in this analysis. In all three phases, reconstructed z positions of data events are lowered by 5 cm to account for a discrepancy between the SNOMAN and the physical detector:

$$z' = z_{\text{FTP}} - 5.0 \quad (6.9)$$

Additionally, in Phase I the mean energy of ${}^{16}\text{N}$ calibration data was observed to slowly decrease over time. This “drift” is explained by changes in detector optics which were not

fully modeled in the first phase, and this is compensated for with a time-dependent energy correction:

$$C(t) = \begin{cases} 1.5902 - 6.3032 \times 10^{-5} \cdot t, & \text{if } t < 9363 \\ 1.1806 - 1.9288 \times 10^{-5} \cdot t, & \text{otherwise} \end{cases} \quad (6.10)$$

where t is the number of days since December 31, 1974 and the corrected energy $T'_{\text{eff}} = T_{\text{RSP}}/C(t)$. For the salt and NCD phases, the simulation was improved so as to track changes in the detector optics and no such correction is required. However, there is a small modification to remove some residual drift in the NCD phase [47]:

$$C(t) = 1.197 - 1.751 \times 10^{-5} \cdot t \quad (6.11)$$

where t is again the number of days since the reference date. A correction is also applied to the β_{14} isotropy parameter for the Phase I Monte Carlo:

$$\beta'_{14} = 0.9919 \times \beta_{14},$$

which accounts for mis-modeling of the Rayleigh scattering in Phase I, based on a comparison to the Phase II model for which the model parameters were explicitly measured [46].

In addition to the observable corrections, adjustments are also required in the normalizations of some signals. These corrections, described fully in Section 8.6 of Reference [27], compensate for events aborted during simulation, modeling inaccuracies, and dead time introduced by cutting bursts of events, in order to better match the acceptance in Monte Carlo to data:

- The standard solar model implemented in SNOMAN is the BP2000 model; ^8B and hep solar neutrino fluxes are scaled to their values in the more recent BS05(OP) model.

- The true electron and deuteron density is higher than in the detector model, by 1.31% and 1.29%, respectively [44]. These factors are used to reweight CC, NC, and ES interactions.
- In order to account for radiative corrections in the NC $\nu - d$ cross section not included in the model, NC event rates are scaled down by 2.3%.
- CC interactions on Na and Cl make a small contribution in Phase II; this is modeled by increasing Phase II CC rates by 0.2%.
- The low-level cuts result in small a loss in live time: 1.4% for Phase I, 1.1% in Phase II, and 1.96% in Phase III, and Monte Carlo is reweighted to account for this.
- Due to events aborted in the simulation, an energy-dependent correction to Monte Carlo event weights is made: $W = 1/(1 - 6.238 \times 10^{-4}T_{\text{eff}})$.

6.8 Counting Analysis Overview

In analogy to the 2006 SNO Phase I *hep/DSNB* search [15], I perform a counting analysis wherein events are subjected to tight high-level cuts, counted, and compared to background expectation in a frequentist (Feldman-Cousins) framework. This yields results simple to interpret and to compare to the existing publication, and can also provide model-independent limits.

Systematics are propagated in the counting analysis via a Monte Carlo method. A large ensemble of pseudo-experiments is constructed with systematics drawn from the appropriate distributions, and a counting analysis is repeated for each. The ensemble is sampled such that correlated parameters are fixed across phases, so that the final sum is effectively an ensemble of full three-phase SNO experiments. In cases where systematics are correlated among themselves, for example in energy scale and resolution measured from fits to

calibration data, linear correlations are accounted for. The resulting joint distributions are used to estimate the uncertainties in the expected signal and background rates, and the median upper limits for the fluctuated ensemble are used to determine the overall sensitivity estimate.

There is an additional systematic uncertainty in the counting analysis associated with the normalization of the ${}^8\text{B}$ rate, which I determine by extrapolating a fit to low energy (6–12 MeV) data into the *hep* energy region of interest. The statistical uncertainties on that fit contribute to the total systematic uncertainty on the background estimate.

The counting analysis is fully described in Chapter 8.

6.9 Signal Extraction Overview

A multi-dimensional signal extraction fit can improve the limit (or measurement uncertainty) beyond the counting experiment for the *hep* flux by leveraging shape information correlated across multiple observable dimensions and correlations in systematic uncertainties.

To this end, I have developed signal extraction code which performs a maximum-likelihood fit using a Markov Chain Monte Carlo (MCMC). This approach uses a random walk to map out the likelihood space in arbitrary dimensions, producing a posterior distribution proportional to the likelihood function which can be used for parameter estimation in either a frequentist or Bayesian framework.

The fits include eight signals:

- *hep* Flux (with CC and ES rates correlated, and the flux correlated across all three phases)

6.9 Signal Extraction Overview

- ${}^8\text{B}$ Flux (with CC and ES rates correlated, and the flux correlated across all three phases)
- Atmospheric neutrino rate (Phase I, II, and III)
- Neutron rate (Phase I, II, and III)

and twelve floating systematics (uncorrelated across phases):

- Energy scale
- Energy scale nonlinearity
- Energy resolution
- Angular resolution (for ES)

in three dimensions:

- Kinetic energy (6–20 MeV)
- Angle to the Sun ($\cos \theta_\odot$)
- Isotropy (β_{14})

Other high-level parameters are applied to the data as cuts prior to the fit, and other systematics are scanned (iteratively shifted and refit). The three-dimensional probability distributions are constructed from SNOMAN Monte Carlo, after corrections have been applied. The fit algorithm requires thorough testing in order to understand any systematic biases or other issues. This is achieved through ensemble testing of Monte Carlo data sets with known levels of signal, including bias and pull testing. The signal extraction analysis is fully described in Chapter 9.

6.10 Physics Interpretation

The physics interpretation of the result is straightforward. The counting and fit-based analysis yield either an upper limit or a measurement of the rate of *hep* solar neutrino and DSNB events observed in some predefined region of interest. There are two ways in which this is used:

Model-independent limits are determined by choosing an analysis region with very low background, and performing a counting experiment to set an upper limit under the assumption that all events are due to signal. This provides an upper limit on the flux independent of, e.g., an oscillation model.

Model-dependent limits, conversely, are determined by comparison to specific model predictions; here, those are the standard solar model (SSM) and the DSNB model proposed by Beacom and Strigari. The signal Monte Carlo is used to convert limits or measurements in terms of an event rate into flux units.

The set of limits (or measurements) for *hep* and DSNB, accounting for all relevant systematic uncertainties constitute the primary results for this analysis. I present both limits in the context of particular solar and DSNB models, as well as raw event rates, such that limits may be derived for arbitrary models.

Chapter 7

Characterization of Energy-Related Systematic Uncertainties

The dominant ${}^8\text{B}$ solar neutrino background for the *hep* search has a steeply falling energy spectrum near the *hep* endpoint, and so the measurement of the *hep* flux is highly sensitive to the detector’s energy response. The SNO energy reconstruction algorithms described in Section 5.6 have been extensively validated through comparisons between data and Monte Carlo simulations [46], using a variety of calibration sources (see Section 5.4.2). The present *hep* and DSNB analysis takes place at higher energies than e.g. ${}^8\text{B}$ solar neutrino flux measurements, however, so it must be confirmed that the existing energy reconstruction routines perform adequately in this regime as well.

By comparing reconstructed event energies for nominally-identical data and Monte Carlo data sets, we can measure the correction required to achieve consistency and the corresponding uncertainty. In typical terminology, we are investigating the uncertainty in the detector

resolution, though ‘resolution’ in the sense of a Gaussian convolution is only part of a more general response function. The uncertainty is defined by our ability to reproduce data with the simulation model, and limited by finite statistics available for the comparison. Ultimately, we are deriving constraints on a parametric model of the deviations in energy response that tell us the extent to which Monte Carlo-derived signal and background PDFs can be distorted while still being consistent with the calibration data.

There is a limited amount of source data available to validate the energy response at high energy, 14–40 MeV. In Phase I, the pT source provided 19.8 MeV γ rays, but this was not deployed in Phases II or III. Michel electrons resulting from muons that stop inside the detector, meanwhile, provide a high-energy distributed calibration source with a well-known spectrum in all three phases, that covers the full analysis energy range. The statistics for this sample are relatively poor, however, with only $\mathcal{O}(100)$ events in each phase.

In this chapter I describe an independent re-analysis of the pT source data and a fit to Michel electron data, using the ^{16}N and pT source calibration as constraints. This provides a measurement of energy systematics that are subsequently used in both the counting and signal extraction analyses. It is also crucial to constrain the possible contribution of non-Gaussian tails, which could introduce additional smearing of the ^8B events producing an excess at the endpoint that fakes a *hep* signal. This possibility appears to have been neglected in previous work, including *hep* studies. Therefore, I will also present here a new analysis of the ^8Li source data — which has a spectrum similar to the ^8B background — to constrain the probability of energy smearing into a flat tail.

7.1 pT Source Constraints

The pT ($^3\text{H}(p, \gamma)^4\text{He}$) source, which delivered 19.8 MeV γ rays, was deployed only once (in Phase I) and at three positions along the z axis (0, -250, and -500 cm). The run list for this data is given in Appendix E. This section presents a reanalysis of the pT source data

using the latest detector Monte Carlo model and improved cuts to better reject background.

The pT source γ rays have no identifying tag, and the source produces copious neutrons due to target and beam impurities, which contaminate the low-energy side of the γ peak and must be removed in analysis. The following cuts are applied to select candidate γ events:

- Low-level cuts, with the NHIT burst cut excluded
- Event reconstructs near the source: $R_s < 125$ cm
- Event is directed away from the source: $\hat{u} \cdot R_s > 0.85$
- Electron-like high-level cuts
 - $0.2 < \beta_{14} < 0.8$
 - $\text{ITR} > 0.55$
 - $P_\phi > 0.0001$ or $P_{2d} \cdot E^4 > 3.60$
 - $10^{-4} < \text{ICT} \leq 1$

The R_s and $\hat{u} \cdot R_s$ cuts are powerful because the neutron background has a larger range and is uncorrelated in direction; these distributions for data at the central position are shown in Figure 7.1. Figure 7.2 shows the effect of all the cuts, relative to the raw spectrum with only data cleaning cuts applied. The same position, $\hat{u} \cdot R_s$, and high-level cuts are applied to both the data and Monte Carlo. The energy spectra are uniformly binned into 150 keV bins from 14 to 30 MeV. After cuts have been applied, I fit the residual neutron peak in data with a Gaussian and subtract the fit function from the spectrum to further reduce this background. The spectra after cuts are fit with a Gaussian function in a restricted range from $-1\sigma \rightarrow 2.5\sigma$, or about 16–30 MeV; an example is shown Figure 7.2(b) for data at the origin. Fits are performed independently for each of the three positions, for data and Monte Carlo, and the differences in mean and standard deviation between data and MC are

7.1 pT Source Constraints

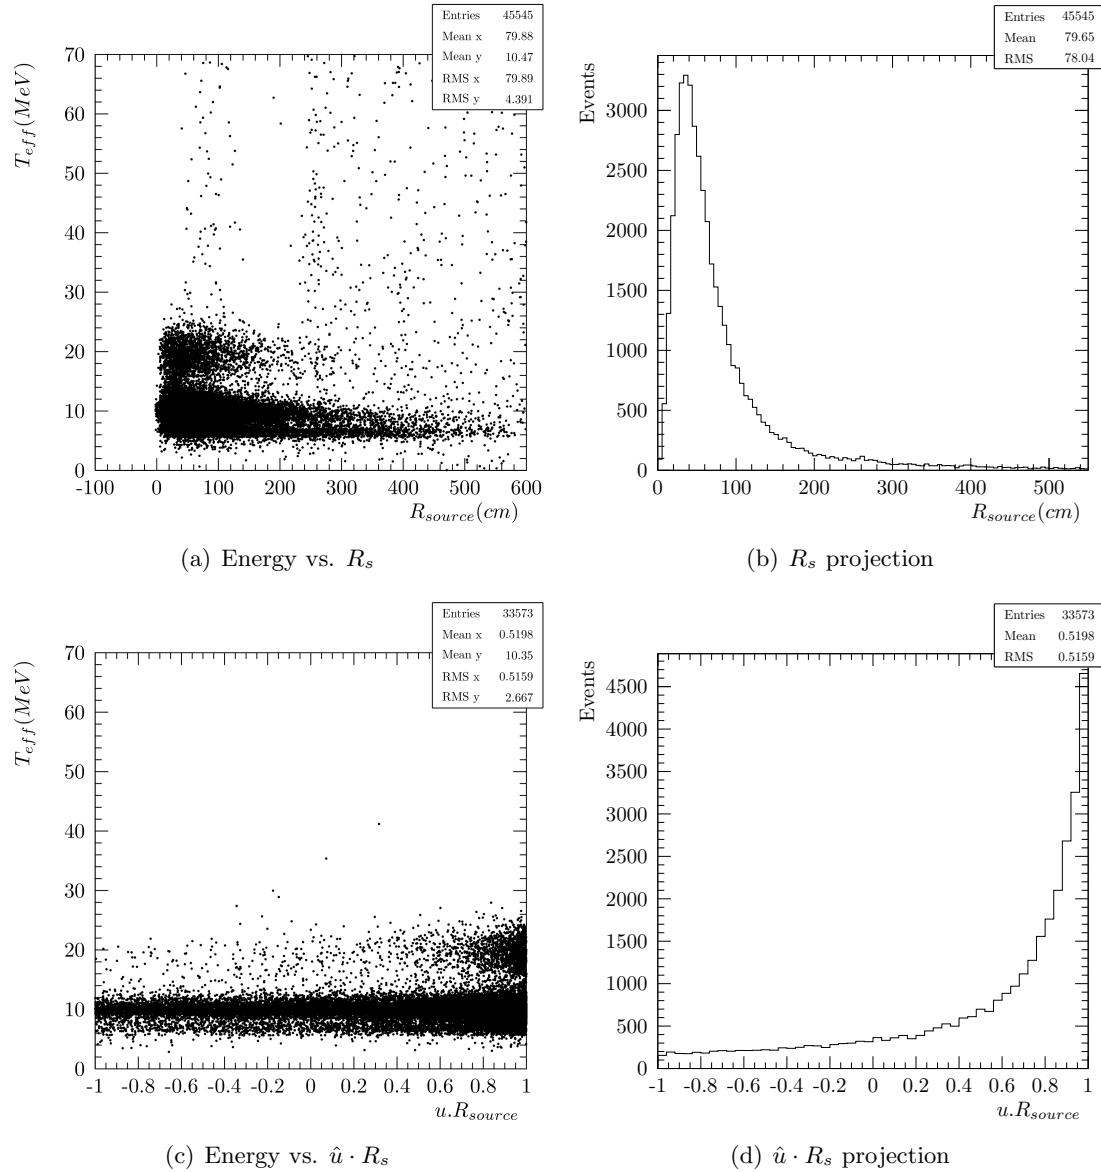
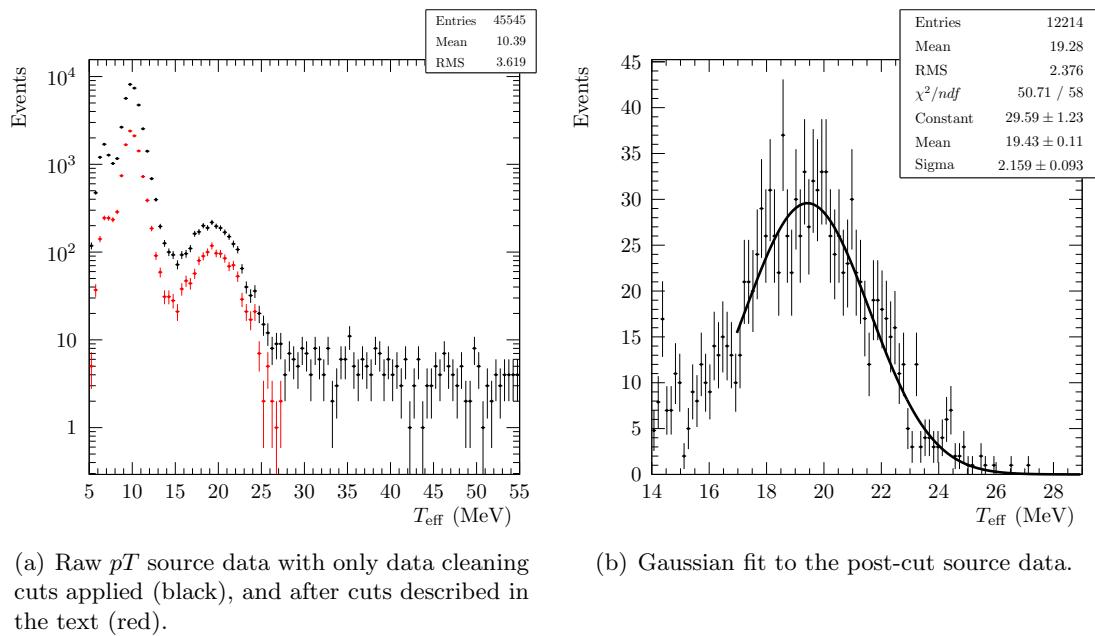


Figure 7.1: Distributions of pT source data at the origin.

7.1 pT Source Constraints



(a) Raw pT source data with only data cleaning cuts applied (black), and after cuts described in the text (red).

(b) Gaussian fit to the post-cut source data.

Figure 7.2: Energy distributions of pT source data at the origin, demonstrating the background reduction due to the cuts and the Gaussian fit to extract the mean and resolution.

taken to characterize the fractional shift in energy scale and energy resolution (i.e. (data - MC) / MC). The best-fit values and shifts are shown in Figure 7.3. It is clear that there is some disagreement between data and Monte Carlo, particularly a radial dependence to the mean energy. Since this source is directional, this could be due either to a real radial bias or a bias in $\hat{u} \cdot r$ intrinsic to the reconstruction algorithm, which would not apply to signal events. Assuming that case we take the most trustworthy point, at the origin, as the central value and take the spread (3.50%) as a systematic uncertainty. This same issue has been observed, and is handled in the same way, in Reference [15].

The resolution shift is consistent with zero, in contrast to the previous analysis where the (volume weighted) average resolution shift was 4.6%. This appears to be due to the improved neutron background rejection, which reduces contamination on the low side of the peak which would tend to inflate the best-fit resolution.

The extracted resolution shift is complicated by the broadening of the γ spectrum relative to monoenergetic electrons, since the former may scatter multiple electrons, or deposit most of their energy in a single scatter. The effect on the resolution is assumed to be normal, in which case the resolution for electrons is related to the resolution of source events by convolution with a Gaussian with width σ_{corr} , such that

$$\sigma_{\beta}^2 = \sigma_{\text{source}}^2 - \sigma_{\text{corr}}^2. \quad (7.1)$$

This broadening correction is determined from Monte Carlo, by comparing the Gaussian fits to isotropic electron MC at $r < 250$ cm. Since MC is generated only at discrete energies, I use the nearest (20 MeV) simulations and scale to the pT energy before performing the fits. Comparing this resolution to the central pT source MC, the extracted correction is $\sigma_{\text{corr}} = 0.842 \pm 0.124$. Systematic uncertainties associated with this correction are neglected, since the result is not used directly but only as input to another fit.

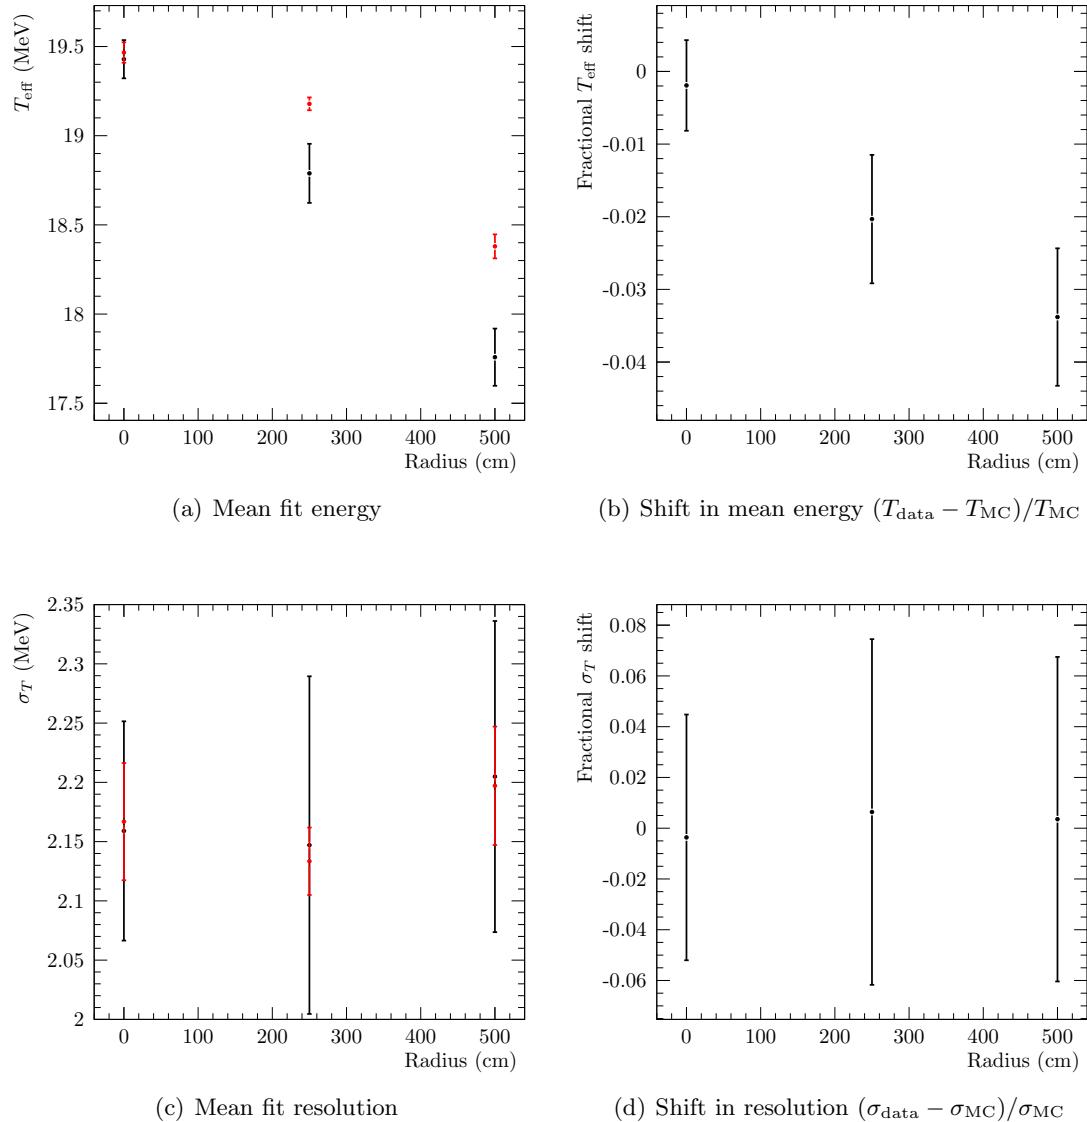


Figure 7.3: pT source data (black) and Monte Carlo (red) comparisons as a function of radius, and the corresponding shifts.

z [cm]	Mean			Resolution		
	Data	MC	Shift	Data	MC	Shift
0	19.43 ± 0.11	19.47 ± 0.06	-0.0019 ± 0.0062	2.16 ± 0.09	2.17 ± 0.05	-0.0036 ± 0.0484
-250	18.79 ± 0.17	19.18 ± 0.04	-0.0203 ± 0.0088	2.15 ± 0.14	2.13 ± 0.03	0.0064 ± 0.0681
-500	17.76 ± 0.16	18.38 ± 0.07	-0.0338 ± 0.0095	2.20 ± 0.13	2.20 ± 0.05	0.0035 ± 0.0639

Table 7.1: Results of pT source fits. Units are MeV, except shifts which are fractional.

Since the resolution fits appear unbiased, I perform a volume-weighted average¹⁴ of those points to derive a single pT resolution. The pT source fit results are given in Table 7.1; the energy shift is $-0.19 \pm 3.38\%$, the source resolution shift is $0.43 \pm 6.51\%$, and the 19.8 MeV electron equivalent resolution shift is $0.57 \pm 7.49\%$.

7.2 Michel Electrons

Decay electrons from stopped μ provide a distributed source with well-known spectrum peaked at around 50 MeV, and a time-based tag can readily provide a pure sample; hence these events are a powerful high-energy calibration source in all three phases.

7.2.1 Event Selection

The Michel electron selection criteria are:

- Fails the retrigger low-level cut (follows another event within 5 μ s)
- Passes the in-time channel cut
- Passes high-level cuts
 - $-0.12 < \beta_{14} < 0.95$

¹⁴ Assuming three bins, from 0–125 cm, 125–375 cm, and 375–550 cm.

τ [μ s]	
Phase I	2.21 ± 0.27
Phase II	1.99 ± 0.16
Phase III	2.25 ± 0.22

Table 7.2: Results of lifetime fits to the selected decay electron events.

- ITR > 0.55
- Passes fiducial volume cut ($r < 550$ cm)
- Energy in the range $10 < T_{\text{eff}} < 70$ MeV
- Separated from the first event in the retrigger burst by $> 0.8 \mu\text{s}$ (to remove electronics noise triggers immediately following large events)

These cuts are designed to select electron-like events that follow another event, using the coincidence with a μ event to tag the decay electron. The selection is validated by fitting the distribution of time differences between the predecessor and retrigger event, which should be an exponential distribution with a decay constant near muon lifetime of $2.2 \mu\text{s}$ ¹⁵ ($t_{1/2} = 1.52 \mu\text{s}$):

$$f(t) = Ae^{-t/\tau}. \quad (7.2)$$

The time distributions are shown in Figure 7.4, and Table 7.2 summarizes the fit results.

These events are then compared with stopped μ Monte Carlo in order to evaluate the performance of the reconstruction. Data and simulations are compared in Figure 7.4, using the *a priori* data and Monte Carlo corrections; the normalization is a nuisance parameter, and Monte Carlo is scaled to match data. Without any further corrections, a χ^2 test for

¹⁵Decreased slightly due to muon capture.

7.2 Michel Electrons

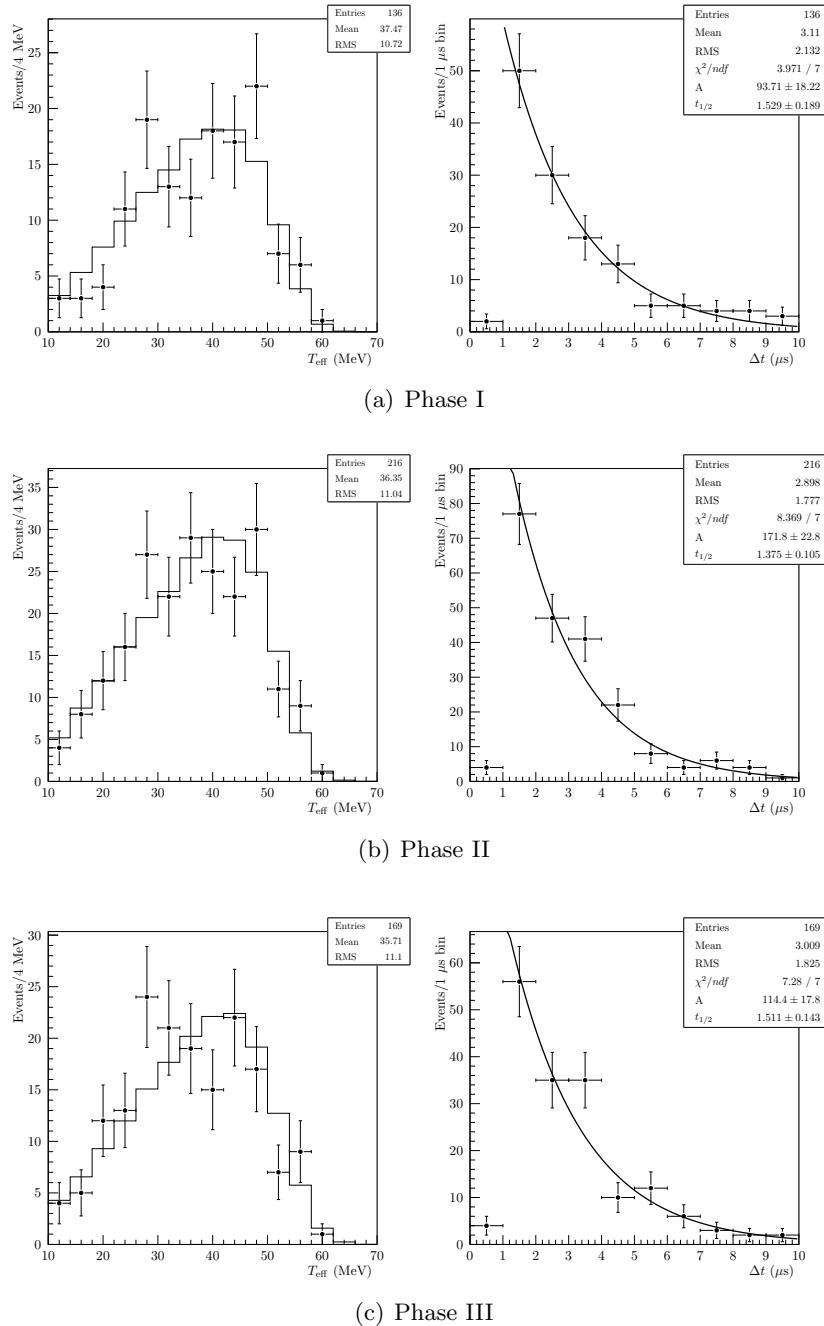


Figure 7.4: Reconstructed energy spectra and time differences between predecessor and re-trigger events for the Michel electron event selection.

these distributions yields a high probability, suggesting that the reconstruction works well at high energy, and systematic corrections should be small.

7.2.2 Energy Fits

I fit the data for the normalization, a fractional shift in the energy scale Δ_S of the form

$$T' = T \cdot (1 + \Delta_S), \quad (7.3)$$

and a fractional change to the energy resolution Δ_R , which pushes the reconstructed energy towards (−) or away from (+) the true value:

$$T' = T + \Delta_R \cdot (T - T_{\text{true}}). \quad (7.4)$$

This parameterization provides some advantages over the commonly-used alternative of a kernel width for a Gaussian convolution (“ σ_{extra} ”). First, the Δ approach allows negative resolution smearing, and although it is unlikely that the Monte Carlo would need to become narrower, asserting the parameter boundary at zero in a fit can result in bias for small values of σ_{extra} . Second, this method is computationally simple, and lends itself well to the signal extraction methods that will be described later. For positive σ_{extra} , the two approaches are equivalent, in that there exists a one-to-one transformation mapping $\sigma_{\text{extra}} \leftrightarrow \Delta_R$ [16]:

$$\Delta_R = \sqrt{1 + \frac{\sigma_{\text{extra}}^2}{\sigma_E^2}} - 1 \quad (7.5)$$

where σ_E is a mean resolution. This mapping will be used to adapt constraints from previous analyses to use in this fit.

The single observable in the fit to Michel electron data is the reconstructed electron-equivalent kinetic energy (T_{eff}), after application of the nominal corrections. The fit optimizes an unbinned extended maximum likelihood function:

$$-\log \mathcal{L} = \tilde{N} - \sum_{i=1}^N \tilde{N} \times P(T_{\text{eff},i} | \Delta_S, \Delta_R) + \mathcal{L}_C \quad (7.6)$$

where N is the number of events in data and \tilde{N} is the normalization parameter, and the systematic parameters determine the shape of the probability distribution P . \mathcal{L}_C contains a set of constraint terms, which come from independent fits of the source data $\Lambda = \{^{16}\text{N}, pT\}$ for Δ_S and Δ_R :

$$\mathcal{L}_C = \sum_{\alpha=S,R} \sum_{\Lambda} \frac{(\Delta_{\alpha}(T_{\text{eff}}^{\lambda}) - \bar{\Delta}_{\alpha}^{\lambda})^2}{2(\bar{\sigma}_{\alpha}^{\lambda})^2} \quad (7.7)$$

where T_{eff}^{λ} is the mean source energy, $\bar{\Delta}_{\alpha}$ is a best-fit value for the independent fit to the source data, and $\bar{\sigma}$ is the source constraint uncertainty.

Minimization of systematic parameters that move events between the finite number of bins is problematic, since the likelihood space is not smooth. I therefore use a Markov Chain Monte Carlo (MCMC) to sample the likelihood space in order to evaluate uncertainties; the MCMC technique is outlined in more detail in Section 9.3.

7.2.3 Parameterization of Systematics

The sources constraints — namely ^{16}N and pT — measure Δ_S and Δ_R at different energies, providing independent constraints on these parameters. In principle, then, the fit is overconstrained, as measurements at three energies are fit with a linear function of energy.

7.3 High-Energy Fit Results

In general, these systematic parameters should be allowed to vary linearly with energy:

$$\Delta_\alpha = \Delta_\alpha^{(0)} + T_{\text{eff}} \times \Delta_\alpha^{(1)}. \quad (7.8)$$

However, it is not clear whether these additional degrees of freedom are indeed necessary. To evaluate this, I performed very loosely-constrained fits with zeroth-order systematic parameters, providing an independent source measurement at ~ 50 MeV. Figure 7.5 shows the results of the source fits; a flat line corresponds to a zeroth-order Δ parameter, i.e. there is a single Δ that is appropriate for all energies, as in Equations 7.3 and 7.4. The resolution scaling is consistent with flat, suggesting that the zeroth-order approximation is sufficient. The energy scale, however, is somewhat better described by a linear fit $\Delta_S^{(0)} = -0.4\%$, $\Delta_S^{(1)} = 0.068\%$ than flat (with a χ^2 about half as probable) indicating a weak preference for including the first-order term.

Based on these findings, and the precedent of including a nonlinear scaling and constant resolution correction in previous SNO analyses, the following parameterization is adopted for the Michel fits, and subsequently applied in the counting and signal extraction analysis:

$$T' = T + (\Delta_S^{(0)} + \Delta_S^{(1)} \cdot T) \cdot T + \Delta_R \cdot (T - T_{\text{true}}). \quad (7.9)$$

7.3 High-Energy Fit Results

The constraints input to the fit are listed in Table 7.4, and results are given in Table 7.3 and Figure 7.6. The ^{16}N constraints are extracted from Reference [48], wherein the following

7.3 High-Energy Fit Results

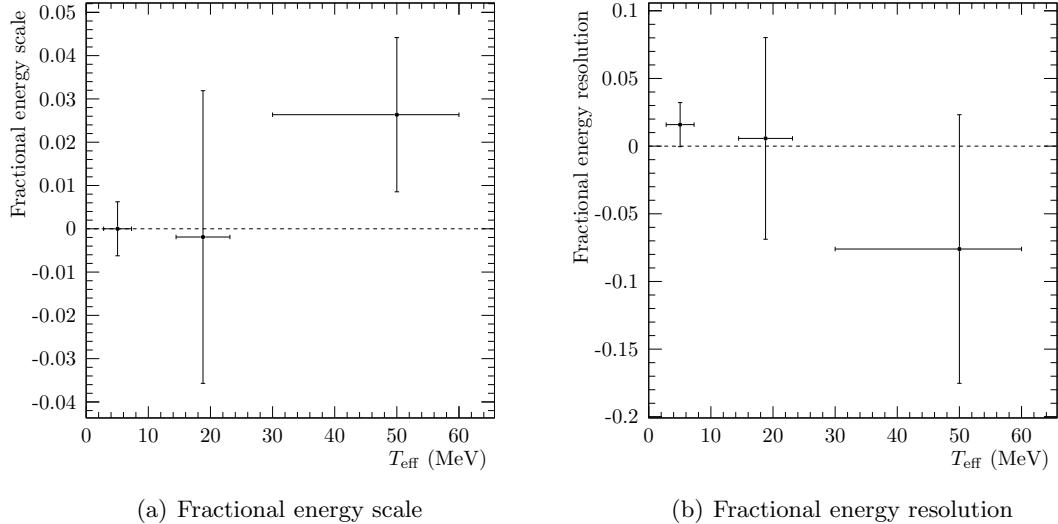


Figure 7.5: Results of independent source fits for energy systematics.

parameterization is used for energy scale and energy nonlinearity:

$$T' = (1 + a_{0c}^E + a_0^E)T$$

$$T'' = \left(1 + a_1^E \frac{T' - 5.05}{19.0 - 5.05}\right) T'.$$

This is trivially transformed into the form of Equation 7.9, taking the ^{16}N energy and propagating the uncertainties on a_i^E . The resolution as implemented in Phase III is already in the same form as used here, but for Phases I and II (LETA) this transformation used the σ_{extra} parameterization defined above, and this is translated to Δ_R using Equation 7.5 and taking [16]

$$\sigma_E = -0.185 + 0.413\sqrt{T} + 0.0254T.$$

7.3 High-Energy Fit Results

Parameter	Phase I	Phase II	Phase III
N	135 ± 12.2	213 ± 14.8	172 ± 13.0
$\Delta_S^{(0)}/10^{-3}$	-5.20 ± 7.21	-0.01 ± 6.14	1.25 ± 10.2
$\Delta_S^{(1)}/10^{-3}$	0.44 ± 0.42	-0.16 ± 0.37	-0.16 ± 0.43
$\Delta_R^{(0)}/10^{-2}$	1.83 ± 1.60	2.38 ± 1.71	1.61 ± 1.37

Table 7.3: Results of the Michel electron fits.

This approach is inconsistent in that the Phase I and II energy systematics in the three-phase analysis are determined using the FTK energy fitter, while this analysis uses RSP. At low energies, however, the differences in the fitters are expected to be small, and these uncertainties are only used indirectly in this analysis as input to the high-energy fit.

The energy scale and energy resolution shifts and uncertainties as a function of energy were evaluated using a Monte Carlo sampling method; the results are shown in Figure 7.7. As noted previously, these parameters and their correlated uncertainties serve as input to both the counting experiment (for systematics sampling) and the signal extraction fit (as a constraint on floating systematic parameters). In both cases linear parameter correlations are captured by using the covariance matrix to model the likelihood space as a multivariate normal distribution.

As a final cross-check on the fit, I compare the results to a simpler approach where the source measurements shown in Figure 7.5 are fit with zeroth- and first-order polynomials; the results for Phase I are shown in Table 7.6. These parameters are consistent with the full Michel fit, with slightly larger uncertainties, as expected.

7.3 High-Energy Fit Results

	Energy [MeV]	$\Delta_S/10^{-2}$		$\Delta_R/10^{-2}$	
		Constraint	Fit	Constraint	Fit
Phase I ^{16}N	5.0	0.0 ± 0.62	-0.30 ± 0.62	1.59 ± 1.63	
Phase I pT	18.8	0.91 ± 3.38	0.31 ± 0.69	2.68 ± 11.3	1.83 ± 1.60
Phase I Michel	50	—	1.70 ± 1.77	—	
Phase II ^{16}N	5.0	0.0 ± 0.53	-0.08 ± 0.53	1.87 ± 1.76	
Phase II Michel	50	—	-0.79 ± 1.61	—	2.38 ± 1.71
Phase III ^{16}N	5.0	0.0 ± 0.91	0.04 ± 0.91	1.61 ± 1.41	
Phase III Michel	50	—	-0.69 ± 1.76	—	1.61 ± 1.37

Table 7.4: Constraints used in the Michel electron energy systematics fit, and the fit results evaluated at the constraint energies.

	N	$\Delta_S^{(0)}$	$\Delta_S^{(1)}$	Δ_R
N	1.000	-0.007	0.008	-0.006
$\Delta_S^{(0)}$	—	1.000	-0.593	0.032
$\Delta_S^{(1)}$	—	—	1.000	-0.015
Δ_R	—	—	—	1.000

Table 7.5: Correlation coefficients for Phase I Michel fit.

Phase I	
$\Delta_S^{(0)}/10^{-3}$	-3.11 ± 7.40
$\Delta_S^{(1)}/10^{-3}$	0.57 ± 0.45
$\Delta_R/10^{-2}$	1.31 ± 1.57

Table 7.6: Results of direct fits to the source data shown in Figure 7.5.

7.3 High-Energy Fit Results

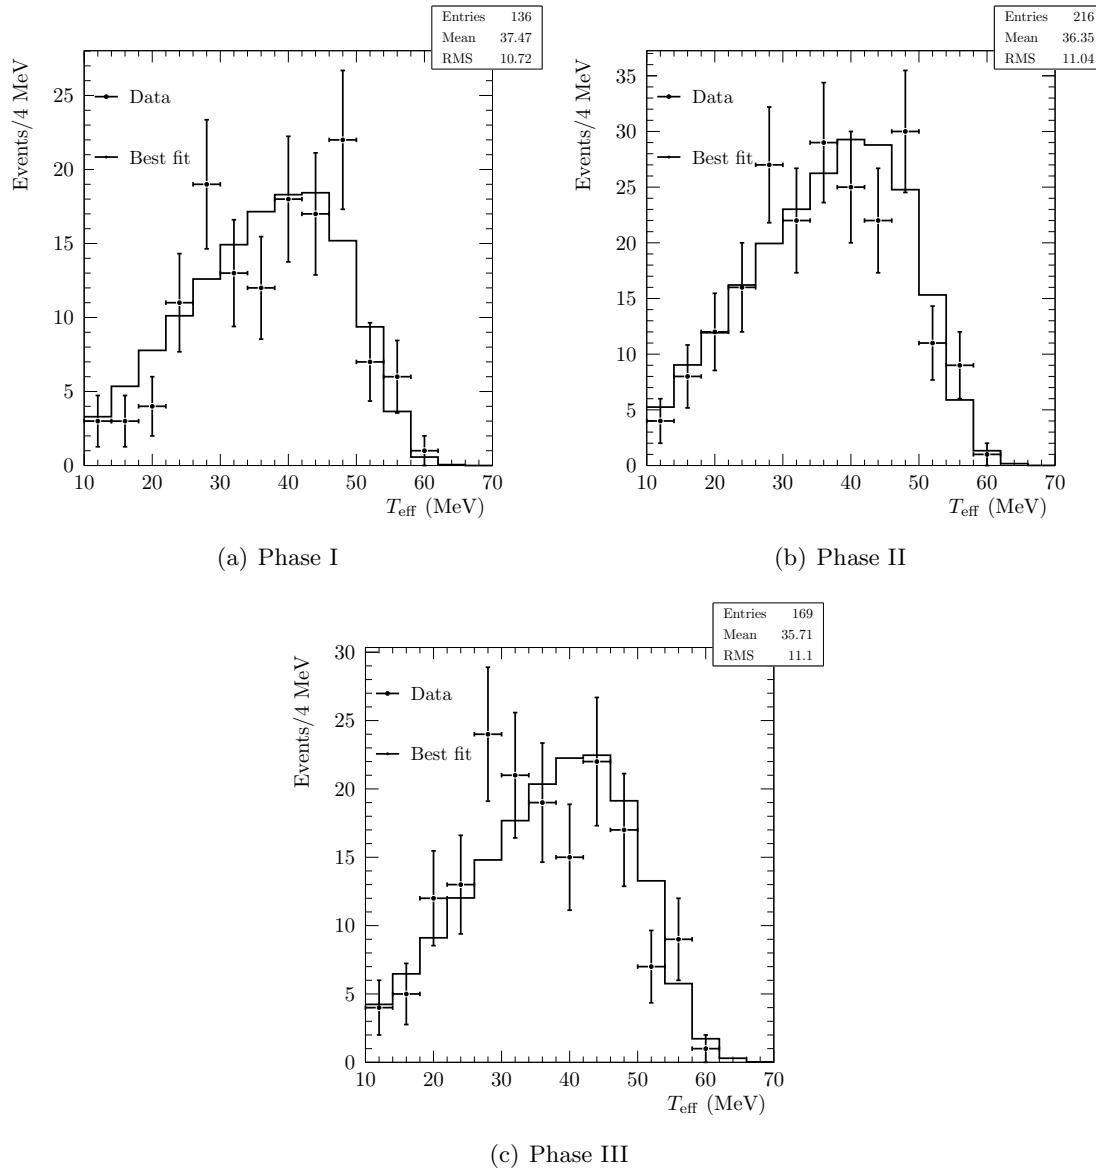


Figure 7.6: Best fit to Michel electrons for energy systematics.

7.3 High-Energy Fit Results

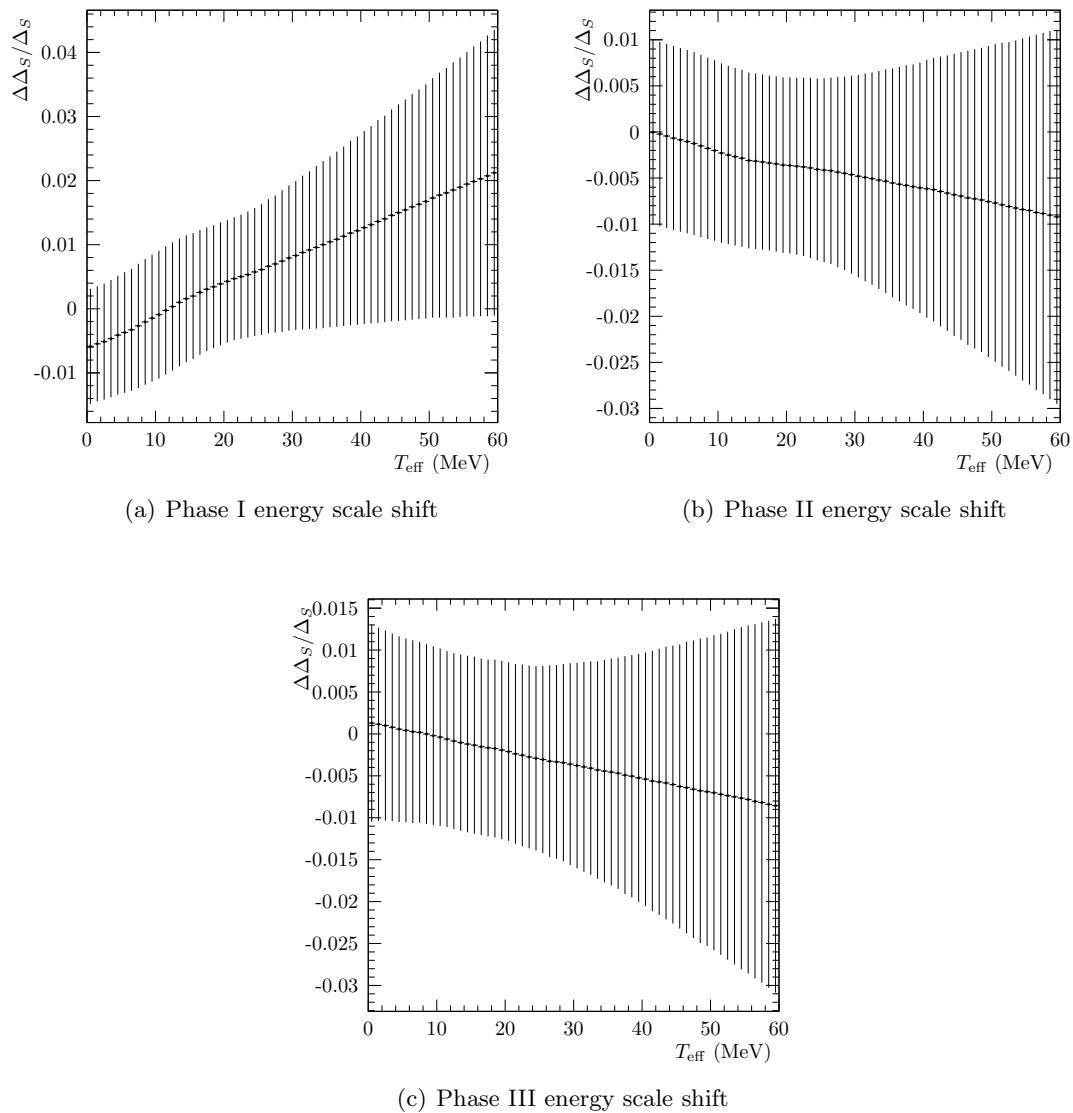


Figure 7.7: Energy dependence of the scale parameters extracted in the Michel electron fit.

7.3.1 Ensemble Testing

In order to trust in the results of the fit and the uncertainties, ensemble tests were performed using a large number of fake data sets sampled from the joint PDF distributions and with rates Poisson-distributed around the expected mean, and constraints sampled within their respective uncertainties.

Figure 7.8 shows the results of the ensemble testing for Phase I, including the parameter pull and log likelihood distributions. The pull is defined, for a best-fit value r and true value \bar{r} , as

$$\text{Pull}(r) = \frac{r - \bar{r}}{\sigma_r}, \quad (7.10)$$

with σ_r being the uncertainty on r reported by the fit. One expects that, for example, the true value lies within the 68% CL fit uncertainty in 68% of the fits. The pull distribution is constructed such that, for a large ensemble of experiments with accurately calculated errors, it tends toward a normal distribution with a mean of zero and unity standard deviation. A non-unit width can indicate a problem with the fit uncertainty. For a brief and informative discussion of normality in bias and pull testing, see Appendix B of Reference [27]. The log likelihood distribution, on the other hand, provides an indication of whether the data is being modeled correctly; the fit NLL should be distributed the same as fake data sampled from around the true value.

The pull distributions for all parameters in all phases are summarized in Figure 7.9.

7.4 Constraining Tails: ${}^8\text{Li}$ Source

In order to make a compelling observation of the *hep* solar neutrino flux, we must exclude the possibility that an excess is due to a tail in the energy response for ${}^8\text{B}$ events.

The ${}^8\text{Li}$ source provided a tagged source of β decays with an energy spectrum similar to that of ${}^8\text{B}$ solar neutrinos, and with Q -value of 16 MeV; these two decays are analogs in the

7.4 Constraining Tails: ${}^8\text{Li}$ Source

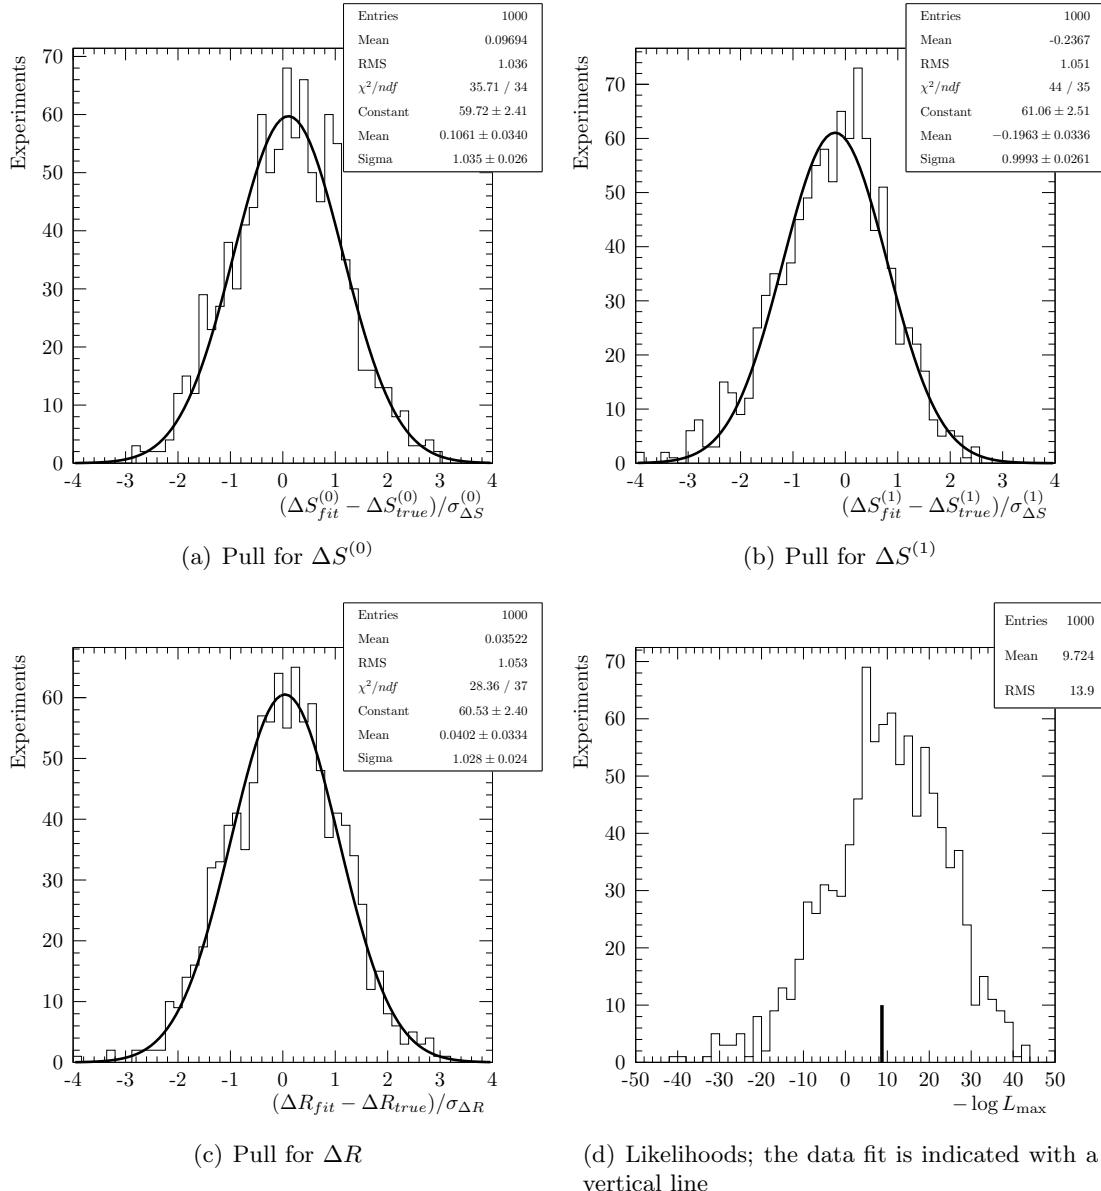


Figure 7.8: Ensemble testing distributions for the Phase I energy systematics fit using Michel electrons.

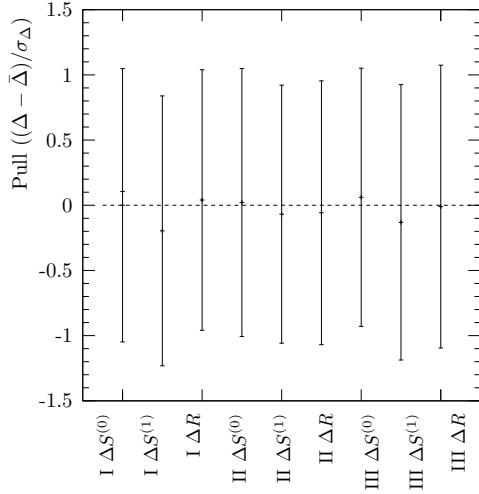


Figure 7.9: Summary of pull distributions for energy-related systematic parameters.

$A = 8$ family [81, 107]. This data was used extensively in SNO for studies of low-level cut sacrifice (see Section 6.6) and position reconstruction calibration, and also for some energy linearity studies. The closeness to the ${}^8\text{B}$ background also means this source is useful for studying departures from the Monte Carlo model of the energy response; the absence of a tail in ${}^8\text{Li}$ data would provide strong evidence that ${}^8\text{B}$ has no tails either.

The source data has contamination from ${}^{16}\text{N}$ decays in the chamber, produced by ${}^{16}\text{O}(n, p){}^{16}\text{N}$ activation of source materials. This background can be reduced substantially with a cut on the integral charge detected in the trigger PMT [81], shown in Figure 7.10(a). A cut is also placed on the distance between the deployed source location and the reconstructed event position, R_s , the distribution of which is shown in Figure 7.10(b).

Source runs for each phase were selected based on run length, detector conditions, and the quality as reported by detector operators in shift reports; the list of selected runs is provided in Appendix F. To select candidate ${}^8\text{Li}$ events, the following cuts were applied:

7.4 Constraining Tails: ${}^8\text{Li}$ Source

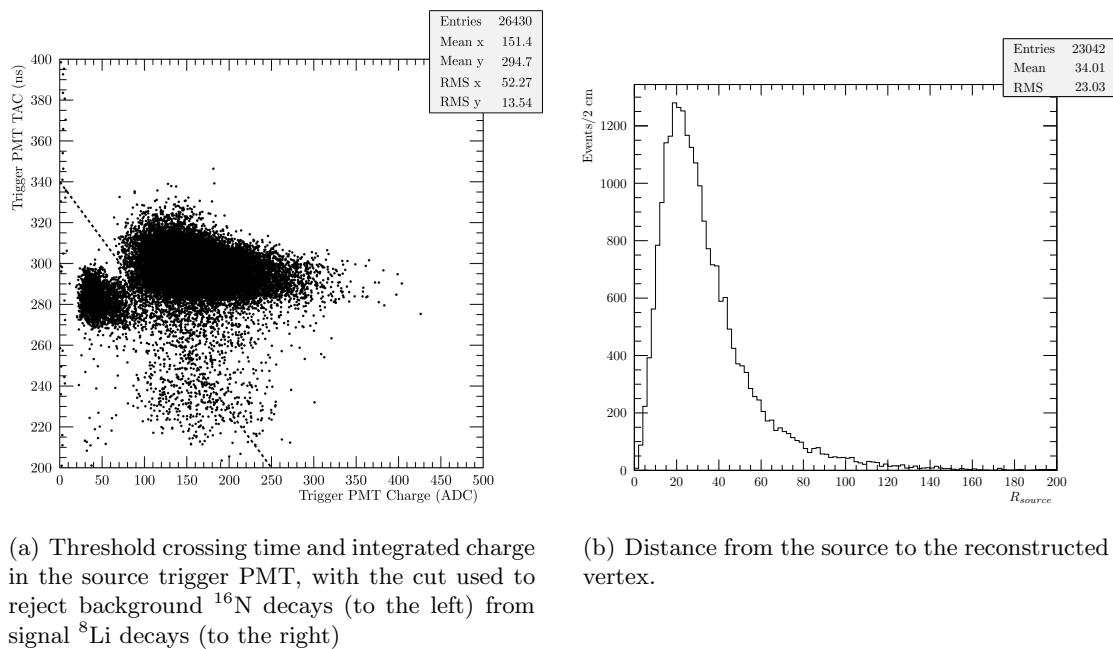


Figure 7.10: Cut parameter distributions for Phase I central ${}^8\text{Li}$ source run 14348.

- The raw detector trigger and the source PMT trigger both fire
- ${}^8\text{Li}$ -like trigger PMT pulse (see Figure 7.10)
- $R_s < 70$ cm
- Reconstructed radius $r < 800$ cm
- Electron-like high-level cuts
 - $-0.12 < \beta_{14} < 0.95$
 - $\text{ITR} > 0.55$

Note that the high-level cuts are as loose as the weakest used in the subsequent analysis, making the tail constraint potentially slightly conservative.

Each source run was individually simulated in SNOMAN, including a model of the source geometry, with comparable statistics to the data itself, and for this analysis I have produced additional, higher-statistics Monte Carlo data (about ten times the data). Figure 7.11 shows a comparison of data and MC after cuts, where I have added together all the runs in each phase in order to improve statistics and sample the detector volume (albeit in a nonuniform way). For this comparison, the MC normalization has been scaled to match the true number of observed events, and the energy scale, nonlinearity, and resolution has been calibrated using ${}^{16}\text{N}$, pT , and Michel electron data as detailed in the above sections. The agreement between data and simulation suggests that the simulation in general provides an excellent model of the data.

The highest-energy event, with an RSP reconstructed kinetic energy of 18.2 MeV, occurs in the Phase II data, in Run 23058. This event is consistent with a source event in all respects, easily passing the tightest high-level cuts with a well-formed electron-like Cherenkov ring, and with a reconstructed position 20 cm from the source. With 136 PMTs hit, the light yield ($N_{\text{hits}}/\text{MeV}$) is in the center of the distribution for RSP at that energy.

7.4 Constraining Tails: ${}^8\text{Li}$ Source

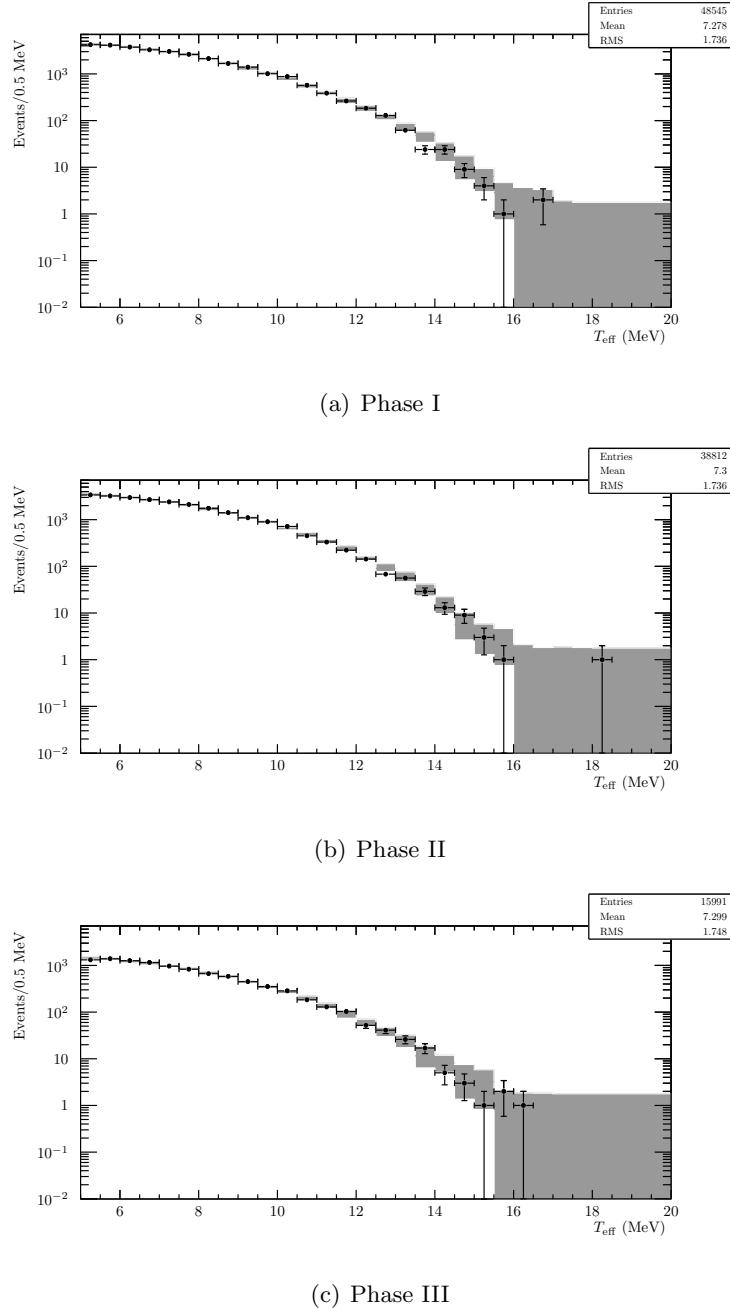


Figure 7.11: Comparison of data (black points) to Monte Carlo (gray bands) for ${}^8\text{Li}$ source runs. The MC bands indicate systematic uncertainties due to energy scale and resolution as well as MC statistics. Source deployment positions are given in Appendix F.

7.4 Constraining Tails: ${}^8\text{Li}$ Source

The probability of pile-up with signal events or low-energy backgrounds (as evaluated using pulsed GT data) after cuts is negligible, and such a high-energy event is not readily explained by radioisotope contamination in the source. The trigger PMT charge and time, on which the ${}^{16}\text{N}/{}^8\text{Li}$ β/α PSD cut is based, is near the maximum density for ${}^8\text{Li}$. Given that the number of hits is large for the ${}^8\text{Li}$ source, and the reconstructed energy is consistent with that N_{hit} , we ascribe this event to a fluctuation (at the $\sim 10\%$ level) in photon production, and not indicative of a problem in the energy reconstruction.

Considering a hypothesis that the energy response has some probability to uniformly distribute events (a flat tail), we may constrain that probability by scanning a likelihood function, $\mathcal{L}(P_{\text{tail}}|\text{data})$, while fixing the other energy systematics, to find $P_{\text{tail}} < 10^{-5}$. Applying this to the ${}^8\text{B}$ background, this is an upper limit of $N_{\text{tail}} \lesssim 10^{-3}$ events in the *hep* search region in each phase, much less than the (approximately flat) atmospheric neutrino background. It follows that this possible tail may be safely neglected, and that the data is adequately described by the model of energy response and systematics at the level of available statistics.

In principle the ${}^8\text{Li}$ data could also provide a high-statistics sample with which to constrain all energy systematics, along the lines of the Michel electron fits. However, this is unfortunately complicated by electron energy loss in the source. In the (EGS4-based) Monte Carlo, ${}^8\text{Li}$ events deposit on average 1 MeV in the stainless steel wall (energy loss in the target volume gas, modeled as air, is negligible). It is difficult to say with confidence whether the sub-percent systematic effects extracted from this data are truly due to energy response, or due to slight mismodeling in the source simulation. Indeed, attempts to perform these fits resulted in an implausible negative resolution shift (i.e. resolution for data is narrower than Monte Carlo), suggesting that this energy loss may be slightly under-predicted. In any case, this data is used only as a general cross-check and to derive constraint on tails.

Chapter 8

Counting Analysis

In analogy to the previously published *hep* and DSNB searches in Phase I [15], I first measure the fluxes using a counting analysis. Within this framework, the events in the dataset that pass an array of cuts (e.g. energy, radius, particle ID parameters, etc.) are counted, and a Monte Carlo model (validated with calibrations) is used to estimate the number of background events expected to pass these same cuts. With that in hand, we evaluate the significance of the observation within the context of the background expectation using standard statistical methods. The following sections detail an analysis along similar lines to the published results, but expanded to include all three phases and making use of the latest detector model, updated high-level cuts, and the energy calibration described in the previous chapter.

I begin by retuning the high-level cuts with which candidate signal events are selected. Next, I describe a signal extraction fit over a restricted energy range (6–12 MeV) used to estimate the ${}^8\text{B}$ background in the *hep* ROI; this is used to reweight the background distributions. Finally, the energy window for the analysis is optimized. The overall process

is as follows:

1. Produce Monte Carlo in SNOMAN for each run to be analyzed
2. Apply corrections (Section 6.7) and low-level cuts (Section 6.3.2.1)
3. Reweight events according to the Monte Carlo oversampling scale factor and normalization corrections (Section 6.7)
4. Optimize and apply high-level cuts (Section 8.1)
5. Perform low-energy signal extraction and reweight events based on the resulting scale factors (Section 8.2)
6. Determine the optimum energy region of interest (Section 8.3)
7. Count surviving events inside the fiducial volume and energy region of interest

In order to account for systematic uncertainties, a Monte Carlo approach is used. I sample many sets of the systematic parameters (including correlations among phases and related parameters) and apply these to create a large ensemble of three-phase pseudo-experiments (i.e. background and signal PDFs), then perform the counting procedure on each. This set is then used to calculate sensitivities (using the background PDFs) and limits (comparing the background estimate to the observed data). The median sensitivity/limit is quoted as the result, and this method also yields the overall systematic uncertainties on the signal and background expectations.

Throughout this chapter, three signal regions are considered, for the *hep* and DSNB searches, and the low-energy signal extraction:

Signal Extraction The range from 6–12 MeV is used to measure the ${}^8\text{B}$ flux, and this result is extrapolated into the *hep* ROI to determine the background rate and uncertainty.

hep The *hep* energy region of interest ranges from about 14 MeV up to 20 MeV; the optimization of the lower edge is deferred to Section 8.3.

DSNB The DSNB energy region of interest spans from 20–40 MeV, with the lower bound also discussed in Section 8.3.

8.1 High-level Cut Optimization

The high-level criteria used to classify events were introduced in Section 6.3.2.2. While in previous work these parameters have been used largely to cut low-energy backgrounds due to radioactive decays (see e.g. References [48] and [46]), here they are used to reduce the atmospheric neutrino background to the single-electron signal. Therefore we must reconsider which high-level parameters are useful, and certainly re-optimize the cut values. Happily, due to the higher energies involved, this generally means tightening the bounds. The cuts represent several orthogonal ways to ask whether an event is compatible with the hypothesis of a single electron. Atmospheric neutrino interactions that produce, for example, multiple Cherenkov rings, short-lived decays that pile up in the same event window, or muons break this hypothesis and these events tend to reconstruct poorly and fail these event quality cuts; Figure 8.1 shows a few examples of such background event classes. Clearly an approach able to interpret multi-ring events would be an improvement, but these naïve cuts are effective for reducing the atmospheric background to a fraction of an event in the 1/3 non-blinded dataset.

Each high-level cut is individually tuned on Monte Carlo to optimize the ratio $S_R/\sqrt{N_R}$ where S_R and N_R are the number of signal events and total events in a region of interest R . This yields 54 independent parameters (6 HLC bounds \times 3 phases \times 3 ROIs), reduced substantially by applying the same parameters in all three phases in most cases, where the values are similar. In each ROI, the dominant signal of interest (i.e. ${}^8\text{B}$ CC, *hep* CC, or

8.1 High-level Cut Optimization

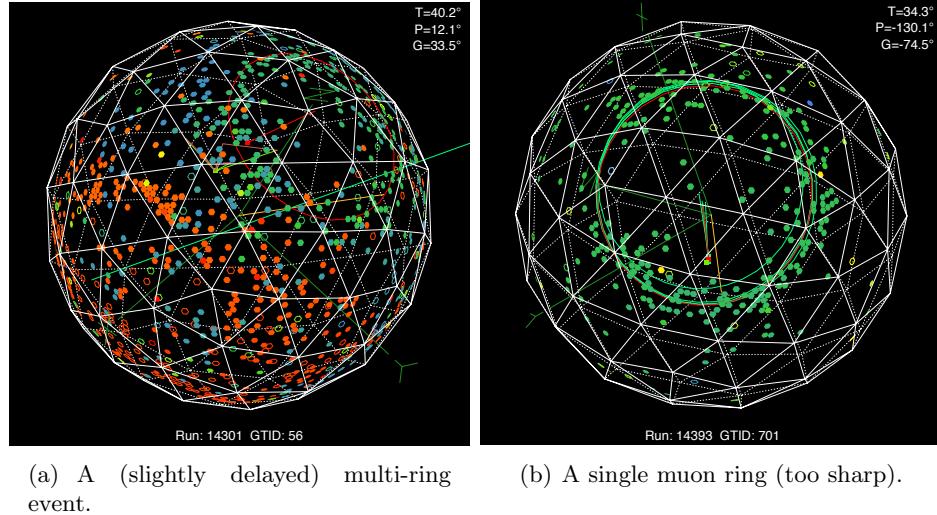


Figure 8.1: Examples of atmospheric neutrino Monte Carlo events removed by the high-level cuts. The color indicates time a PMT was hit: green is earlier, red is later.

DSNB CC) is compared to atmospheric neutrino Monte Carlo with only burst cuts applied. As an example, Figure 8.2 shows the data and MC distributions of the P_ϕ parameter, and 8.3 shows the scans of all parameters for the *hep* ROI in Phase I. In principle, there may exist nontrivial correlations between the high-level parameters, in which case a composite variable could provide better discrimination. However, preliminary tests with linear (Fisher) and nonlinear (boosted decision tree) multivariate approaches suggest that the cuts are largely orthogonal and the potential gain is small.

The approach described above is entirely based on Monte Carlo, and so we must verify that the MC accurately reproduces data. Biases in either the signal sacrifice or the background rejection as a function of energy may affect our interpretation of the data and the extracted limits. Validation of the background rejection is challenging since the frequency of these events in data is quite low, and there exists no source data to use as a proxy. As a general cross-check, however, I consider events with $20 < T_{\text{eff}} < 100$ MeV (assumed to be

8.1 High-level Cut Optimization

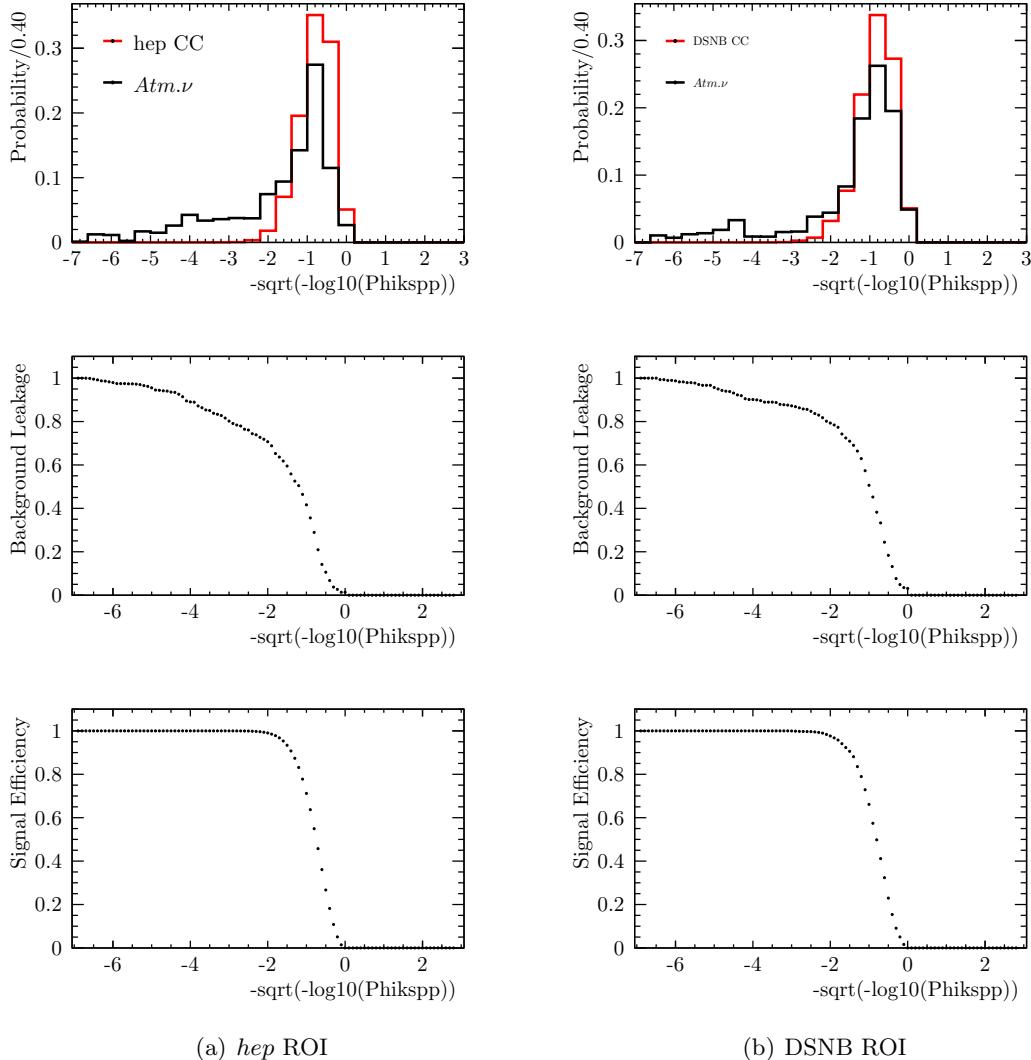


Figure 8.2: Data and Monte Carlo distributions for the 1D angular KS test (P_ϕ) high-level cut parameter in Phase I. The middle and bottom panels show the leakage fraction of atmospheric neutrino background and the sacrifice of the signal as a function of cut value.

8.1 High-level Cut Optimization

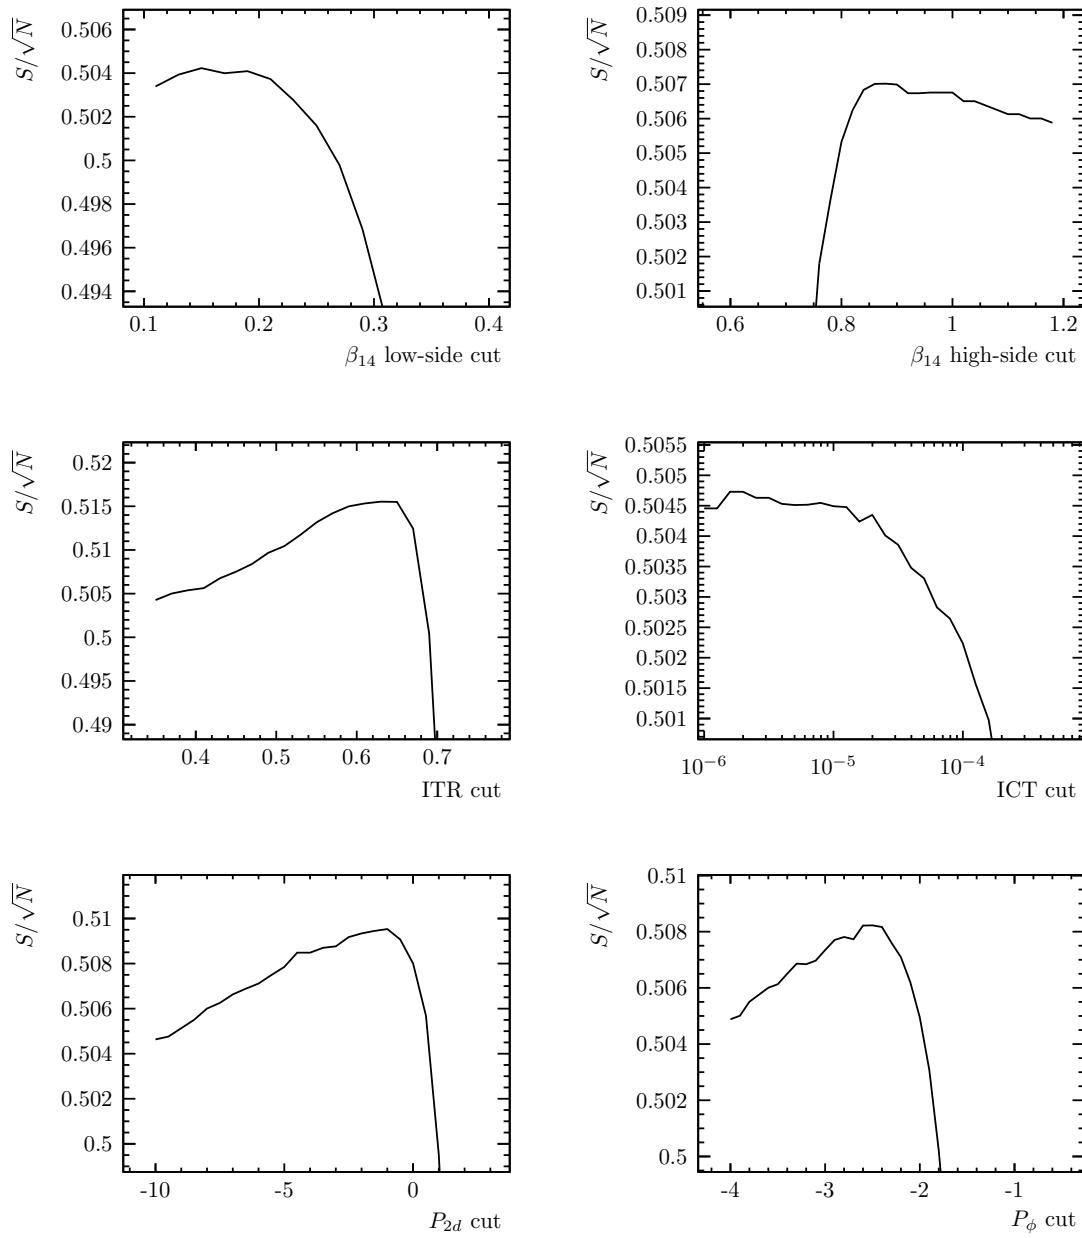


Figure 8.3: High-level cut optimization scans for the *hep* ROI in Phase I.

all atmospheric neutrino-related) that pass the low-level and burst cuts. Across all three phases, there are 17 such events in the 1/3 non-blinded dataset, three of which pass the high-level cuts. Applying the same cuts to the Monte Carlo, about 50% of events are rejected. While this is not a precise measurement of the background rejection efficiency of the high-level cuts, the fact that the data and Monte Carlo are generally compatible provides some evidence that the absolute uncertainty on this small background is itself small, and that we are justified in neglecting the associated systematic uncertainties.

8.1.1 Signal Acceptance

To understand the signal acceptance and ensure that the cuts applied to data have the same effect as on the Monte Carlo (on which the sensitivity estimates are based), I compare data and MC for high-energy calibration sources: the ${}^8\text{Li}$ source, Michel electrons, and in Phase I, the pT source. Figure 8.4(a) shows this comparison for Phase II, using cuts resulting from the MC-based optimization, and it is apparent that there is a significant discrepancy, at least for events following high-energy precursors. This is almost entirely due to disagreement in the ITR distribution, shown in Figure 8.4(b). This is likely due to essentially time-uncorrelated late hits associated with the precursor, such as PMT afterpulses, falling into the follower event window and flattening out the time distribution; the precursor events are not simulated in the decay electron Monte Carlo. To handle this and recover an accurate background estimate, the cut is relaxed to the value used in previous SNO publications, $\text{ITR} > 0.55$. With the ITR cut adjusted, there remains a very large ($\sim 50\%$) sacrifice at high energies; this is due to the ICT cut, which is also relaxed slightly from its nominal value with a negligible loss in sensitivity.

Comparisons of calibration data and Monte Carlo with the final cut values are shown for each Phase in Figure 8.5. Although statistics for high-energy sources are sparse, the distributions are consistent, and so we may expect that the post-cut Monte Carlo distributions

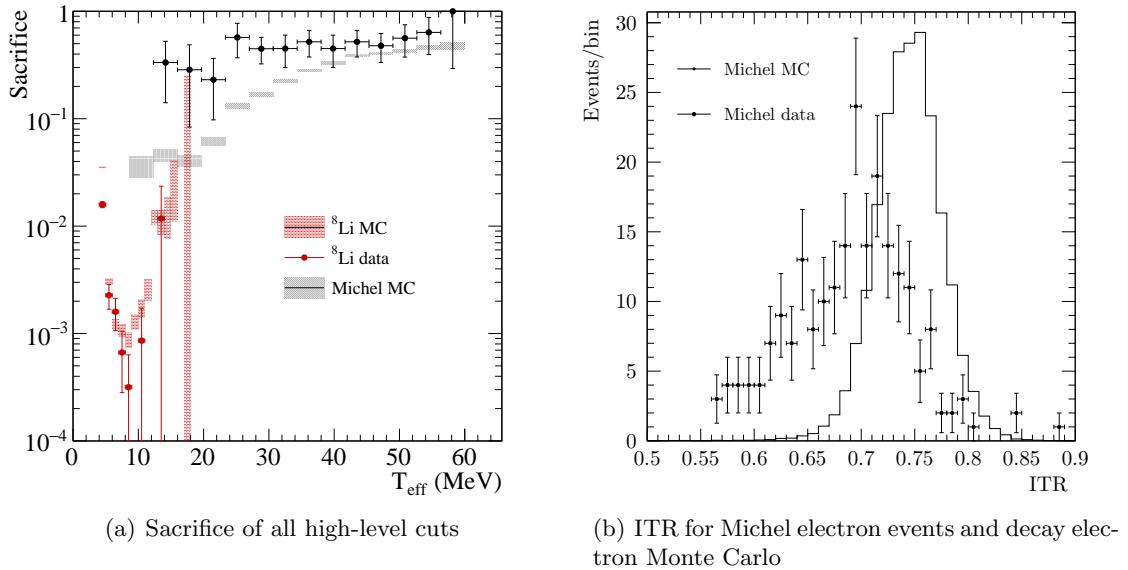


Figure 8.4: Discrepancy between the data and MC for the sacrifice of the level-cuts due to ITR. Phase II data is shown.

used in the sensitivity analysis faithfully represent the data.

To gain a general idea of how the various cuts contribute to the total sacrifice, I have used simulated isotropic electron data as a proxy for the (single-electron) signals of interest with high statistics across all relevant energies. The total sacrifice and that for each individual high-level cut are shown as a function of energy in Figure 8.6.

There is additionally some signal sacrifice associated with the low-level cuts. This effect and the associated systematic uncertainties are discussed in Section 6.6.

8.1.2 Summary of Cut Parameters

The optimized values for the high-level cuts are given in Table 8.1. The tuning of the final cut for the counting analysis, that on energy, depends on the normalization of the ^8B solar

8.1 High-level Cut Optimization

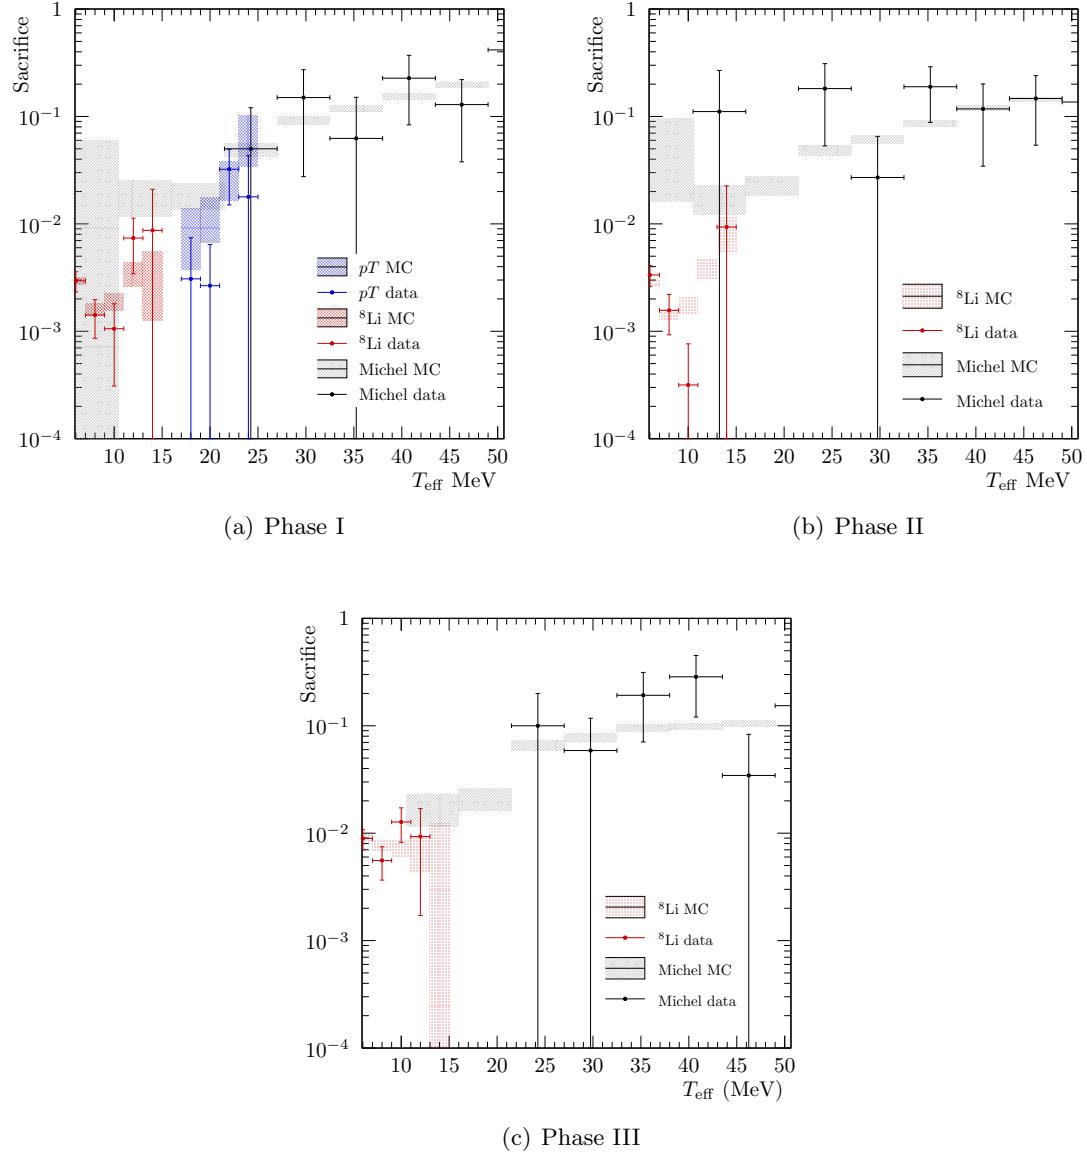


Figure 8.5: Data/Monte Carlo comparison for the total signal sacrifice of the high-level cuts in each Phase.

8.1 High-level Cut Optimization

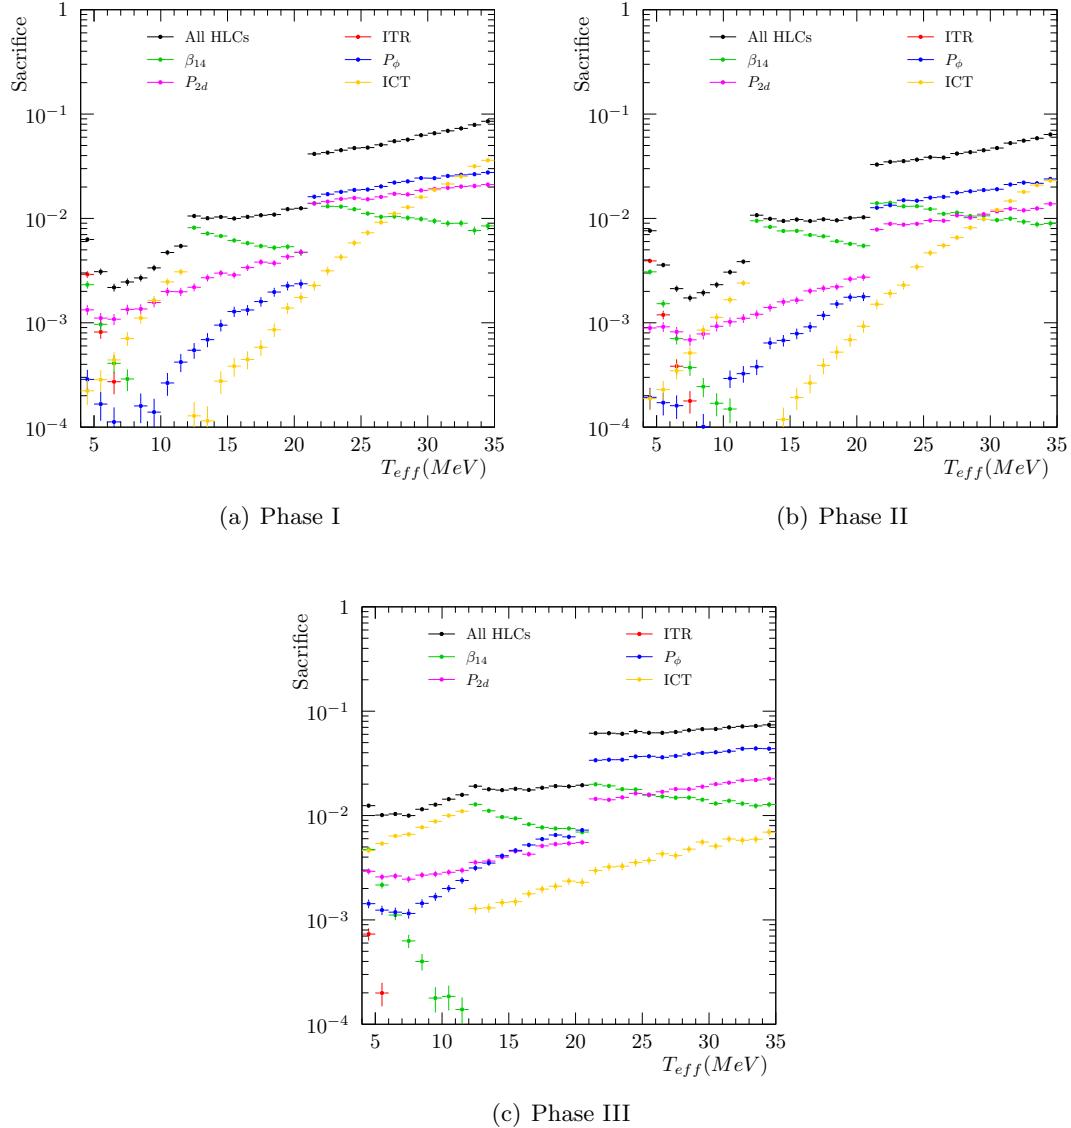


Figure 8.6: Sacrifice in Monte Carlo electrons due to high-level cuts. Discontinuities are due to the three different ROIs which in general have different cuts applied.

8.2 Low-energy Signal Extraction

Parameter	Sig. Ex. ROI	<i>hep</i> ROI	DSNB ROI
Fiducial volume [cm]	550	550	550
β_{14}	-0.12 – 0.95	0.20 – 0.82	0.28 – 0.82
ITR	> 0.55	> 0.55	> 0.55
ICT (Phases I, II)	> 10^{-4}	> 10^{-6}	> 10^{-6}
ICT (Phases III)	> 10^{-2}	> 10^{-3}	> 10^{-3}
$P_{2d} \cdot E^4$	> 10^{-1}	> 10^{-1}	> 1
P_ϕ	> 10^{-6}	> 10^{-6}	> 10^{-2}

Table 8.1: High-level cuts for each region of interest.

neutrino background via the low-energy signal extraction, and is discussed in Section 8.3.

Table 8.2 shows the impact of the burst and high-level cuts on the atmospheric background, while Figure 8.7 shows the final spectrum for major background classes. Together, the burst and high-level cuts reduce the atmospheric background in the *hep* ROI by a factor of 30, and in the DSNB ROI by about a factor of 70, leaving a fraction of an event in both cases, while preserving $\sim 99\%$ and $> 90\%$ of the signal, respectively.

8.2 Low-energy Signal Extraction

The dominant background for the *hep* search is the steeply-falling tail of the ${}^8\text{B}$ solar neutrino spectrum. In order to estimate the number of ${}^8\text{B}$ solar neutrino events inside the *hep* ROI determine the associated uncertainty, I have determined the normalization from lower-energy data in a signal extraction fit. The fit was performed in a restricted energy range from 6–12 MeV, so as to avoid low-energy radioactive backgrounds on the low side and to minimize the contribution of any *hep* signal on the high side. Near 12 MeV, the SSM prediction for the *hep* rate is about 2 orders of magnitude below the expected rate for electrons from the ${}^8\text{B}$ background. The fit includes electrons due to ${}^8\text{B}$ interactions

8.2 Low-energy Signal Extraction

Cut	NCQE		CCQE		All
	15.1 MeV γ	$E_\nu > 100$ MeV	Other		
No cuts	38	616	2100	2754	
β_{14}	33	508	1668	2209	
ITR	33	569	1566	2168	
P_{2d}	33	400	1638	2071	
P_ϕ	35	435	1778	2248	
ICT	36	606	1894	2536	
Pass all HLC	25	366	1243	1634	
Retrigger	38	112	496	646	
Missed μ follower	38	516	828	1382	
NHIT burst	36	539	1676	2251	
Clean event burst	25	387	1077	1489	
Pass all burst	25	67	317	409	
Pass all cuts	17	25	56	98	

Table 8.2: Raw numbers events in the Phase I high-energy atmospheric MC with $14 < T_{\text{eff}} < 20$ MeV and $r < 550$ cm passing each cut. Note that as these numbers are for the raw MC production, and initial state neutrino flavor ratios are not correct.

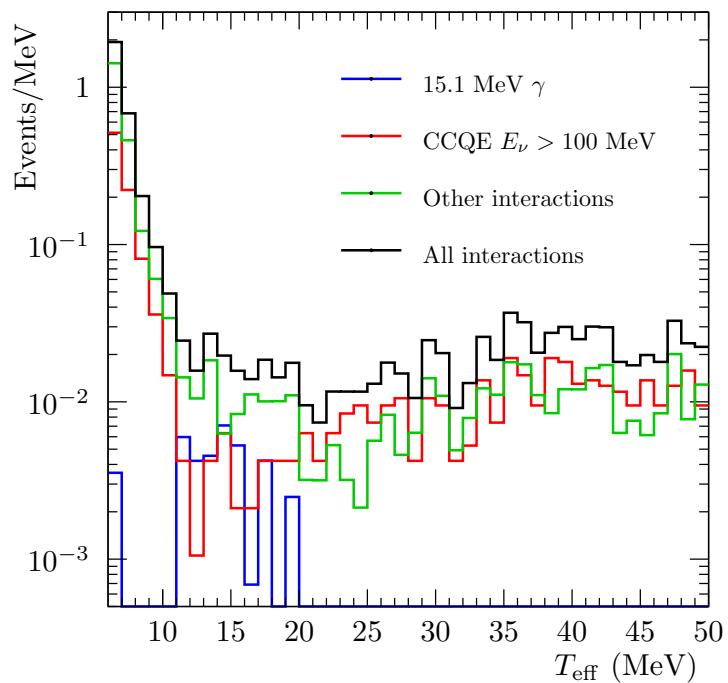


Figure 8.7: Effective electron-equivalent kinetic energy spectra for atmospheric neutrino backgrounds for Phase I after application of the cuts summarized in Tables 8.2 and 8.1.

8.2 Low-energy Signal Extraction

(combining CC interactions on deuterium, ^{17}O , and ^{18}O , and ES on electrons), and neutrons (with a PDF built from ^8B NC events), which are treated as independent. Signal extraction fits are performed separately for each phase.

The signal extraction process consists of an unbinned extended maximum likelihood fit, with optimization performed by MINUIT and error estimation by the MINOS routine [108] as implemented in ROOT 5 [109]. The observables in the fit include the electron-equivalent effective reconstructed energy (T_{eff}), the event isotropy parameter β_{14} , and the cosine of the angle between the reconstructed direction and the direction to the sun, $\cos \theta_{\odot}$. This fit improves on the approach taken in Reference [15] by using fully three-dimensional PDFs (in contrast to using a 2D PDF in T_{eff} and β_{14} and a set of 1D PDFs for $\cos \theta_{\odot}$). The parameters of the fit are the normalizations of the ^8B CC+ES signal and the neutron signal, with PDFs created from SNOMAN Monte Carlo and applying a best-fit three-neutrino oscillation model and all data/Monte Carlo corrections. As this analysis is itself meant to determine a systematic uncertainty, the systematics associated with the mixing parameters are handled separately. The fit allows penalty terms to include *a priori* Gaussian constraints on the parameters. The minimization function is the negative log likelihood:

$$-\log \mathcal{L} = \sum_{j=1}^M N_j - \sum_{i=1}^N \log \left(\sum_{j=1}^M N_j P_j(\{T_{\text{eff}}, \beta_{14}, \cos \theta_{\odot}\}_i) \right) + \sum_k \frac{(N_k - \hat{N}_k)^2}{2\sigma_k^2} \quad (8.1)$$

where N_i and P_i are the mean normalizations and probability distribution, respectively, for each of M signals. In the final term, σ_i is a Gaussian uncertainty on the expected value of the i th signal normalization \hat{N}_i . In practice, the constraint term is not included in the fit.

Fit results are shown in Figures 8.8 (Phase I), 8.9 (Phase II), and 8.10 (Phase III), and summarized in Table 8.3. The results are used to rescale all Monte Carlo PDFs contributing to the ^8B CC and ES and total neutron signals. The results of the Phase I signal extraction

	Phase I	Phase II	Phase III
${}^8\text{B}$ Electrons	$0.969^{+0.051}_{-0.050}$	$0.867^{+0.047}_{-0.046}$	$0.845^{+0.051}_{-0.050}$
Neutrons	$1.229^{+0.260}_{-0.251}$	$1.486^{+0.067}_{-0.066}$	$1.050^{+0.177}_{-0.172}$

Table 8.3: Results of the signal extraction fits with 1/3 data, in terms of the fractional scaling from the model prediction (using BS05(OP) fluxes).

are statistically compatible¹⁶ with the combined oscillation analysis of all three phases published in 2013 [48].

8.2.1 Ensemble Testing

To verify the signal extraction fit and the uncertainties, which will set the expectation for the ${}^8\text{B}$ background in the *hep* ROI and the associated systematic uncertainty, ensemble tests were performed using a large number of fake data sets with rates Poisson-distributed around the expected means.

The pull distributions for Phase I are shown in Figure 8.11. The bias is negligible, but the pull distributions appear too narrow. This is a consequence of the strong correlation between the fit parameters, apparent in fit correlation coefficients (~ 0.54) and visible in Figure 8.11(d). MINOS single-parameter uncertainties are estimated as the extrema of the n -dimensional contour, which is an overestimate for any particular value of the other parameter; this results in a narrowing of the pull distribution. This is not problematic, but symptomatic of using a one-dimensional metric to evaluate a non-trivially multidimensional parameter. This suggests two possible courses of action for the extrapolation of the signal extraction into the *hep* ROI: to use the independent parameter uncertainties, conservatively

¹⁶This signal extraction is performed on the 1/3 unblind dataset, and so statistical compatibility with previous results is all that is expected.

8.2 Low-energy Signal Extraction

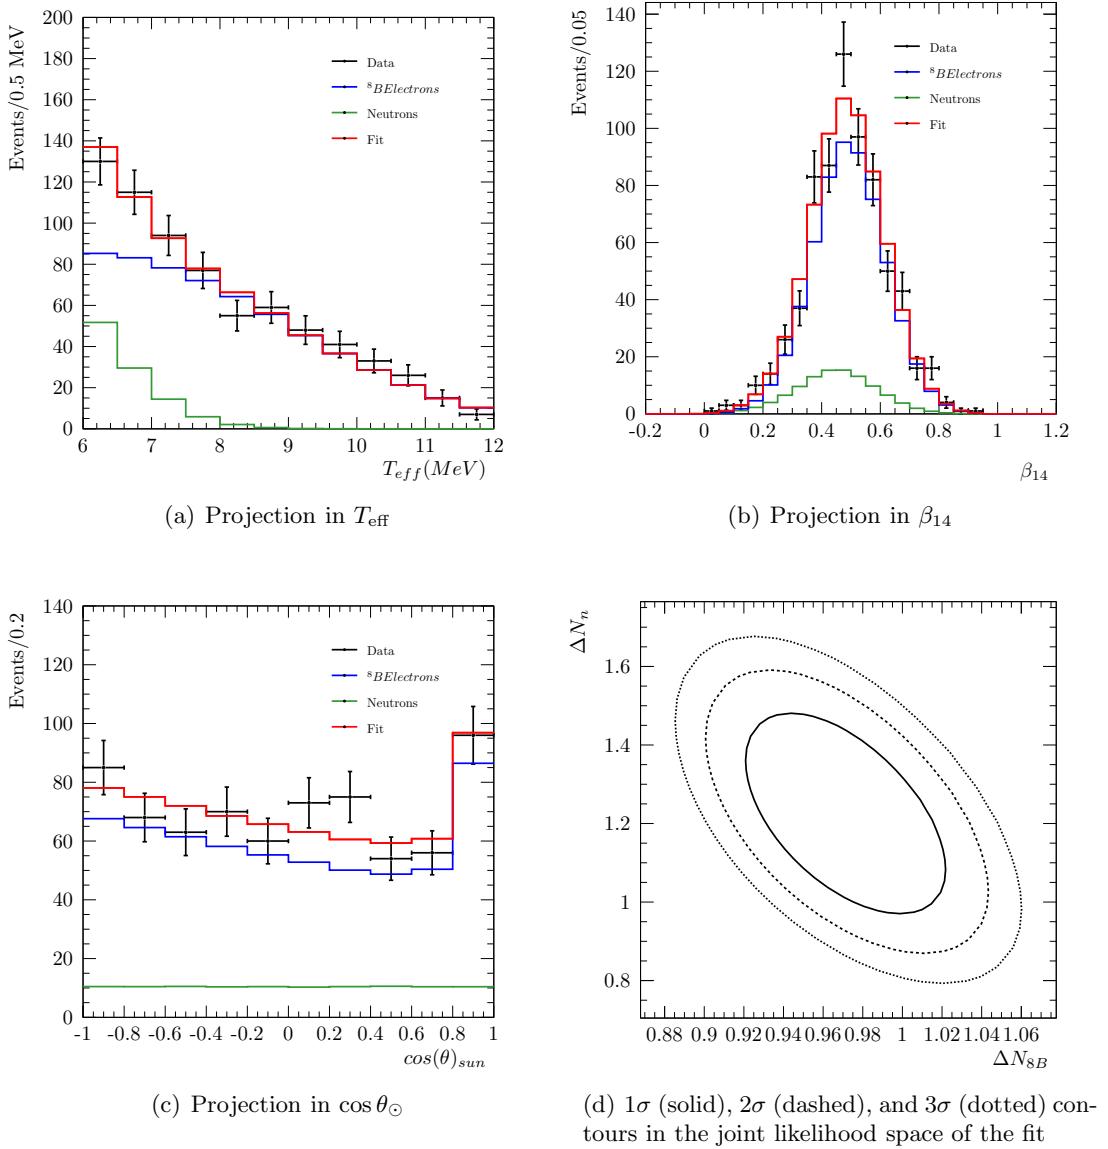


Figure 8.8: Results of the Phase I signal extraction fit with 1/3 data.

8.2 Low-energy Signal Extraction

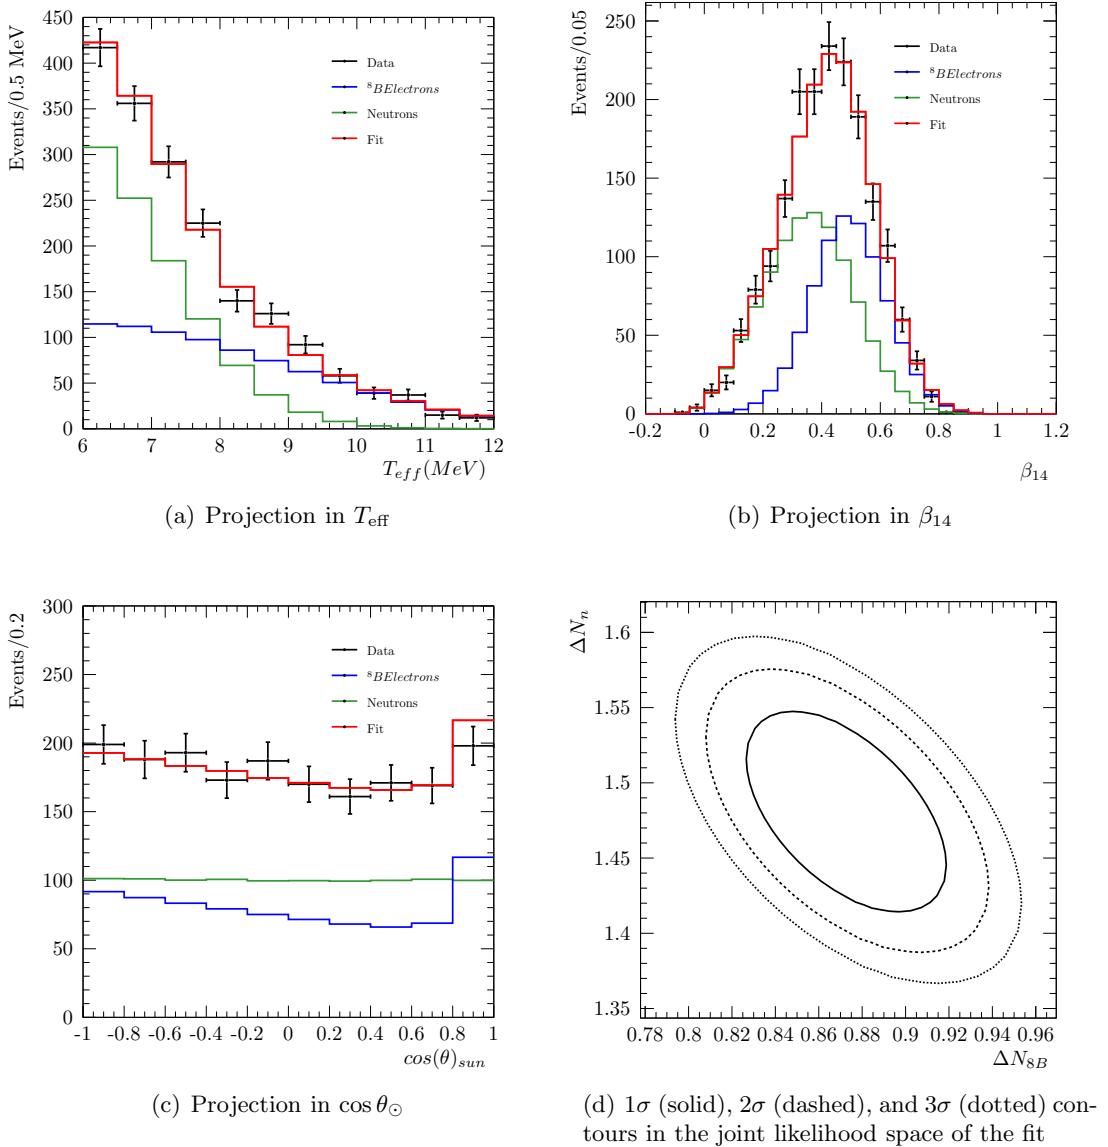


Figure 8.9: Results of the Phase II signal extraction fit with 1/3 data.

8.2 Low-energy Signal Extraction

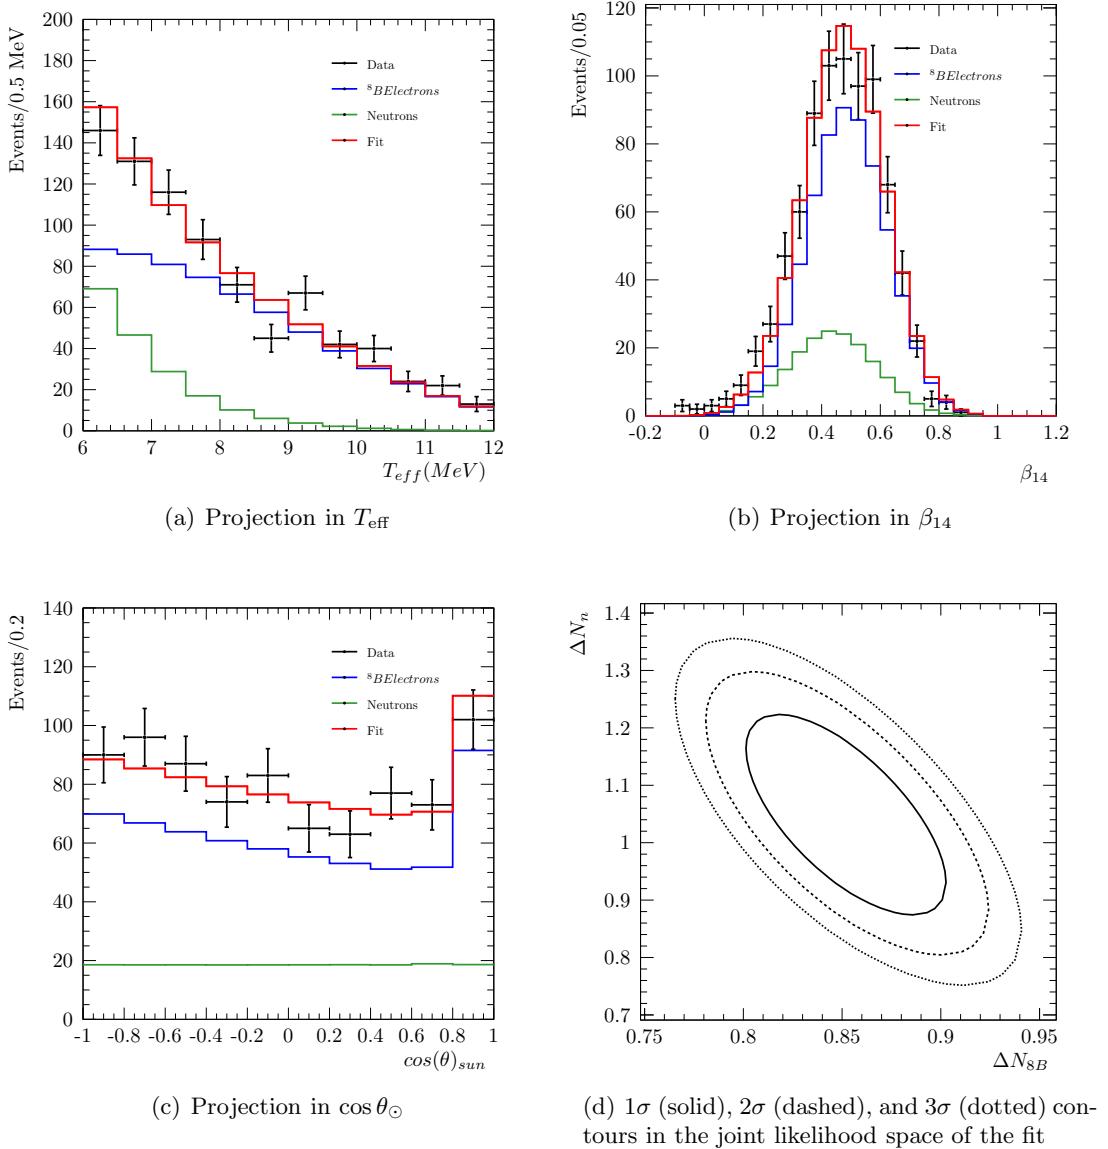


Figure 8.10: Results of the Phase III signal extraction fit with 1/3 data.

treating them as uncorrelated, or to sample from the joint when evaluating systematics in the counting experiment, which will generate correct coverage. The latter approach is taken in this analysis. The pull distributions for both parameters in all phases are summarized in Figure 8.12.

8.3 Energy Window Selection

The final cut made for the purposes of the counting experiment is the energy window “region of interest” (ROI), which is chosen in order to maximize the sensitivity. Using the Monte Carlo data after all corrections and cuts, the expected mean number of signal and background events is calculated as a function of energy threshold. Then, for each threshold, the sensitivity is determined by taking the median upper limit of a large number of sample experiments using the Feldman-Cousins approach [110]. Given the shapes of the ${}^8\text{B}$, *hep*, DSNB, and atmospheric neutrino signals, the optimization is much more sensitive to the low-energy cut than the high-energy cut. The high-side cuts are set at 20 and 40 MeV for the *hep* and DSNB windows, respectively, regions where the sensitivity is essentially flat. The optimization of the lower ROI threshold is performed independently for each phase of the experiment.

Figure 8.13 shows the results of the *hep* and DSNB ROI optimization analysis for Phase I, Figure 8.14 for Phase II, and Figure 8.15 for Phase III. The *hep* sensitivity, shown in terms of a scale factor to the BS05(OP) SSM prediction, is approximately flat up to about 14 MeV. Therefore, while the optimum occurs around 13.5 MeV for Phases I and II, I choose a 14 MeV threshold, improving the signal-to-background ratio with a minimal loss in sensitivity. For Phase III, 14.3 MeV is chosen, owing to the slightly poorer energy resolution in that configuration. For all three phases, the DSNB ROI is chosen as 20 – 40 MeV; this is close to the maximum-sensitivity window, and chosen to be remain insensitive to the *hep* background even if it happens to be on the high end of the BS05(OP) SSM

8.3 Energy Window Selection

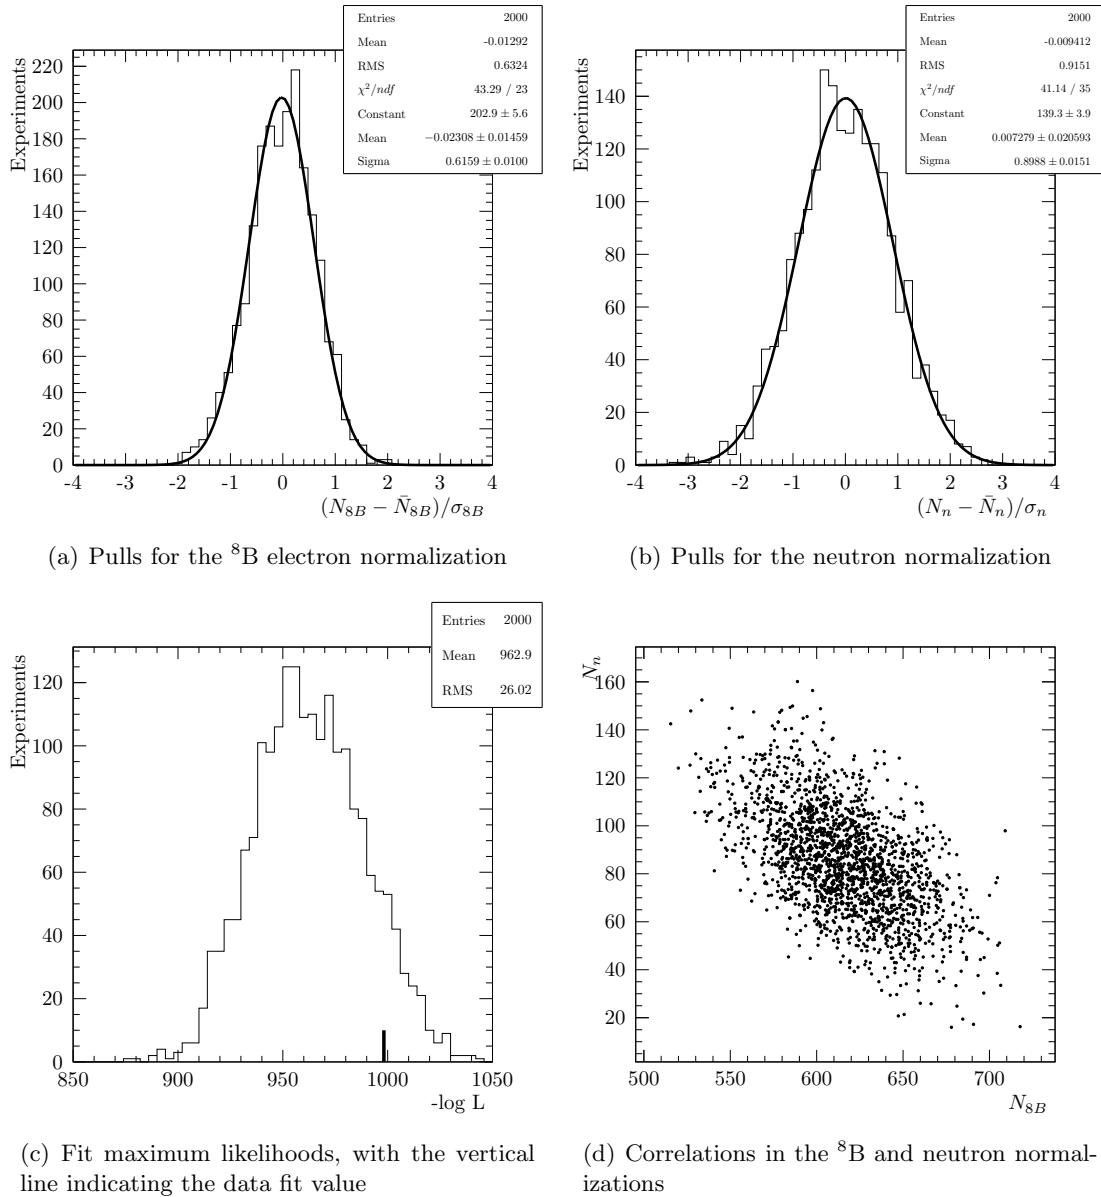


Figure 8.11: Ensemble tests for the Phase I low-energy signal extraction.

8.3 Energy Window Selection

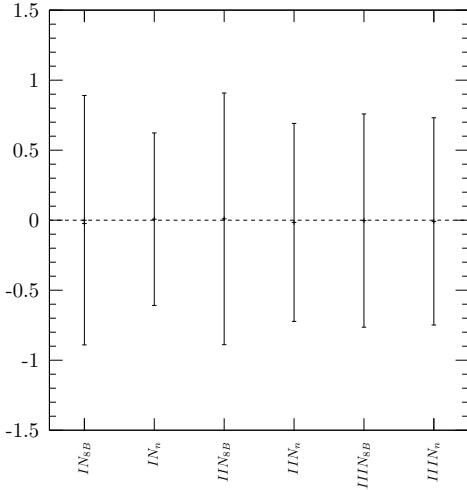


Figure 8.12: Summary of pull distributions for low-energy signal extraction parameters.

model uncertainty. This is a wider window than was used in the Phase I publication (21 – 35 MeV), which was restricted on the high end by certain blindness requirements and on the low end by concerns regarding the then-larger uncertainties on the *hep* flux.

At some level the optimum energy window depends on the systematics, and the selection has been made using the mean values for these parameters. This effect will be small, since the dominant systematics are related to energy response and scale the signal and background together to first order. Nevertheless, a check was performed by varying the threshold ± 0.3 MeV and repeating the full counting analysis, and the above conclusions were found to hold.

8.3 Energy Window Selection

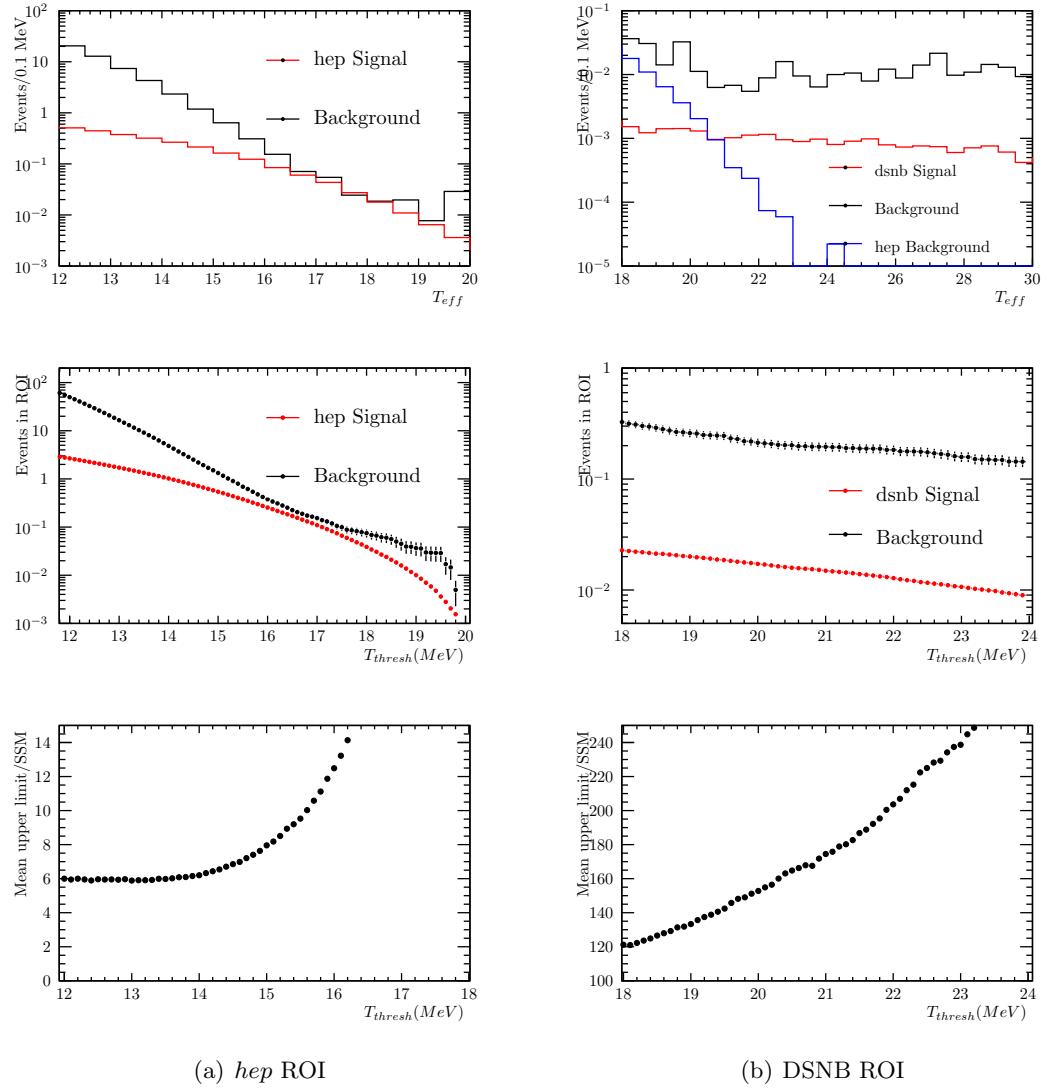


Figure 8.13: The ROI optimization for the Phase I.

8.3 Energy Window Selection

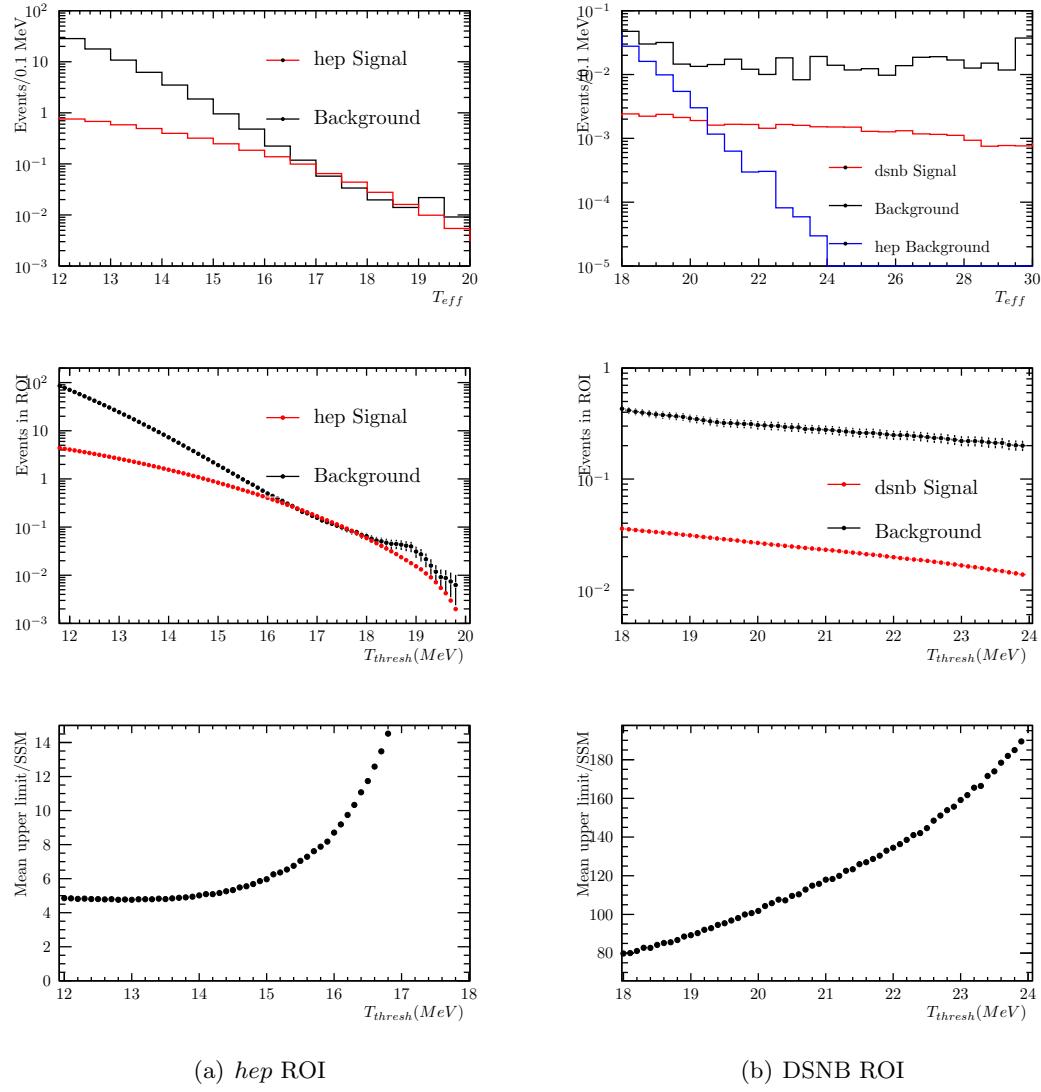


Figure 8.14: The ROI optimization for the Phase II.

8.3 Energy Window Selection

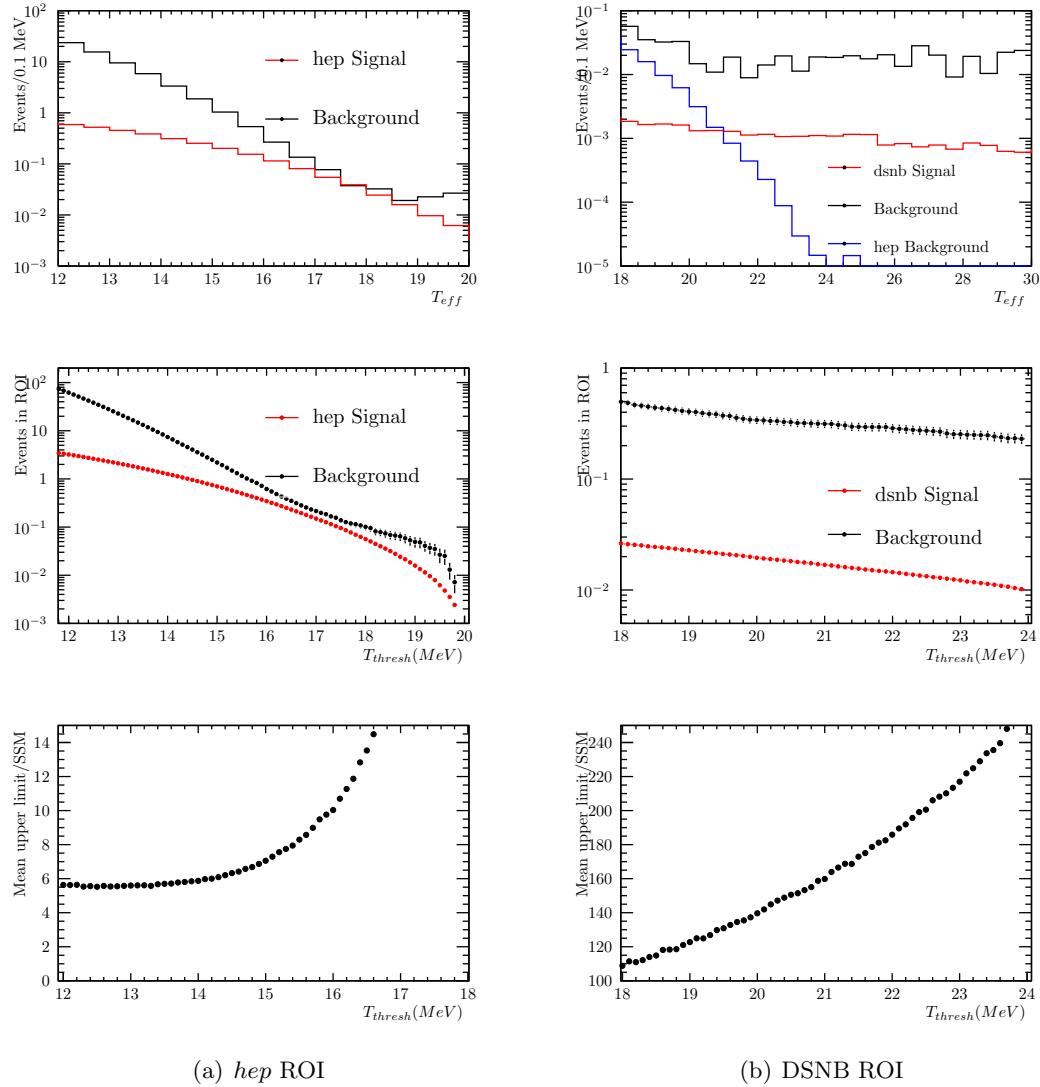


Figure 8.15: The ROI optimization for the Phase III.

8.4 Sensitivity

The sensitivity is calculated using the Monte Carlo signal and background estimate and the Feldman-Cousins framework [110]. The sensitivity provides a statistical metric for the range of parameter values that an experiment should, on average, be able to observe at a given confidence level. The fundamental question of sensitivity is this: what is the minimum signal rate that is greater than a background fluctuation of a given significance? Colloquially, what is the smallest signal that would be visible on top of background fluctuations? Of course, any such estimate must assume that the background model is a complete description of the physical system that generates the data. A brief introduction to the concepts of the frequentist statistical interpretation is provided in Appendix G.

To illustrate the process of sensitivity estimation, the signal and background expectation values for the 1/3 data set are derived with all systematic parameters set to their central values. Note that these are not necessarily the most probable values, but serve to demonstrate the method; these values are not used in the final result. The energy spectra after all corrections and cuts are shown in Figure 8.16, and the contributions to the total background are given in Table 8.4.

For the full counting analysis, this process is repeated many times, for an ensemble of pseudo-experiments where the systematic parameters are sampled randomly (though are correlated, if applicable). This modifies the signal and background PDFs, and in turn the number of signal and background events expected inside the signal region of interest, which has been fixed *a priori*. In this way, we can map a set of PDFs with systematic uncertainties into a set of Feldman-Cousins sensitivities and data-based confidence intervals.

8.4 Sensitivity

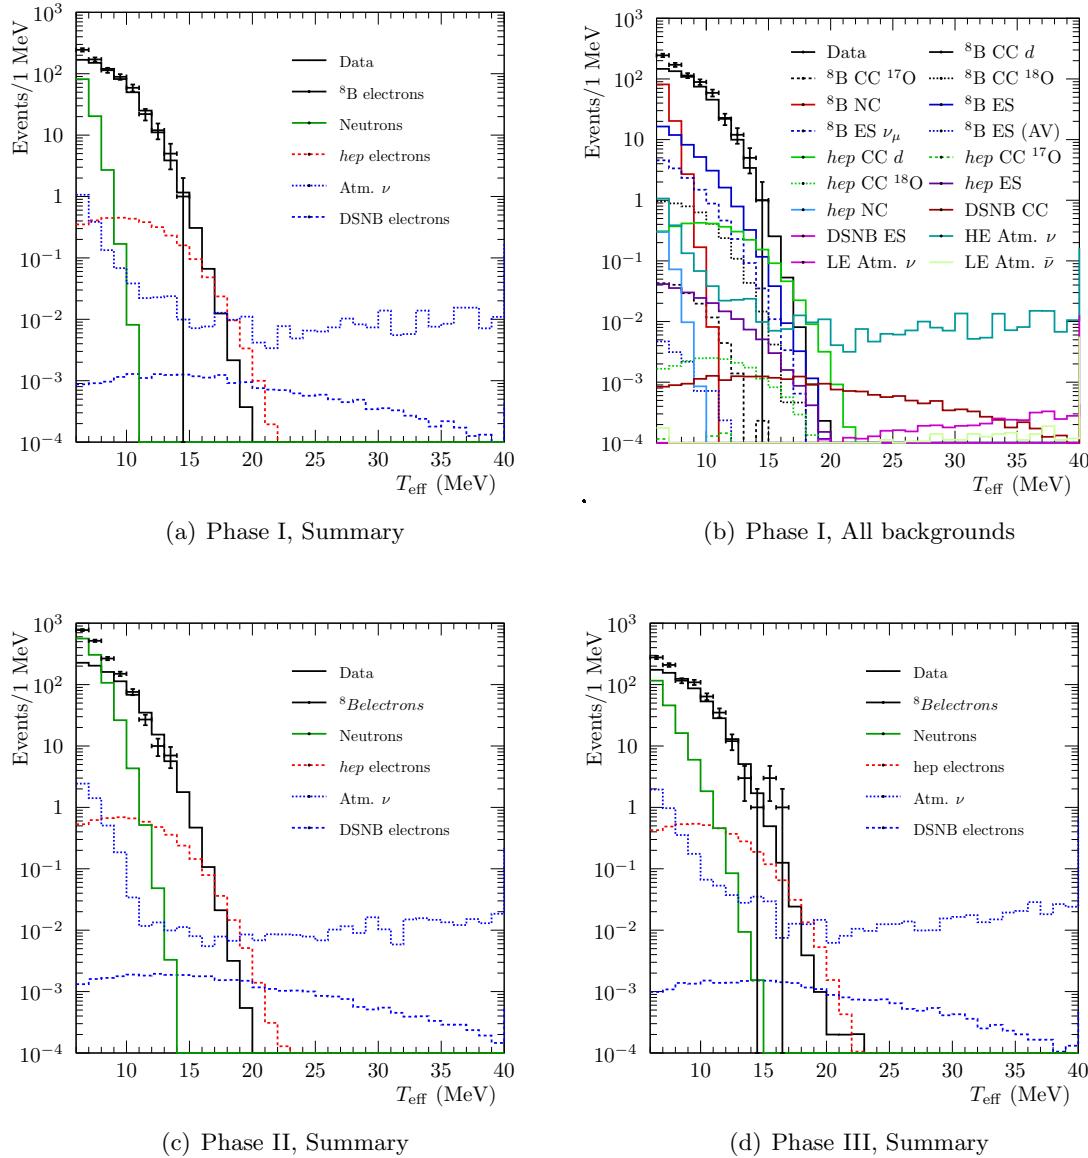


Figure 8.16: Energy spectra around the *hep* and *DSNB* regions of interest, after corrections and all cuts except on energy, with all parameters fixed to their mean values.

8.4 Sensitivity

Signal	Phase I		Phase II		Phase III	
	hep	DSNB	hep	DSNB	hep	DSNB
^8B CC d	0.861095	0.000000	1.351067	0.000000	1.370448	0.000401
^8B CC ^{17}O	0.000000	0.000000	0.000000	0.000000	0.000400	0.000000
^8B CC ^{18}O	0.013454	0.000000	0.024592	0.000000	0.018719	0.000000
^8B NC	0.000000	0.000000	0.000000	0.000000	0.000769	0.000000
^8B ES (ν_e)	0.116836	0.000000	0.178898	0.000000	0.165484	0.000000
^8B ES (ν_μ)	0.035915	0.000000	0.040535	0.000000	0.055347	0.000000
^8B ES AV	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000
hep CC d	0.272934	0.000226	0.416806	0.000420	0.342556	0.000505
hep CC ^{17}O	0.000100	0.000000	0.000166	0.000000	0.000109	0.000000
hep CC ^{18}O	0.002577	0.000000	0.003745	0.000015	0.002941	0.000005
hep ES	0.009358	0.000022	0.014371	0.000032	0.011980	0.000042
hep NC	0.000000	0.000000	0.000000	0.000000	0.000000	0.000000
DSNB CC	0.006208	0.004973	0.009403	0.007673	0.007570	0.005612
DSNB ES	0.000023	0.000003	0.000034	0.000013	0.000037	0.000013
Atm. $\nu E_\nu > 100$ MeV	0.054093	0.059869	0.042063	0.084578	0.098085	0.098260
Atm. $\nu_e E_\nu < 100$ MeV	0.000400	0.001405	0.000730	0.002356	0.000593	0.001825
Atm. $\bar{\nu}_e E_\nu < 100$ MeV	0.000367	0.000888	0.000078	0.000184	0.000919	0.002591
Exp. signal	0.284970	0.004976	0.435087	0.007686	0.357587	0.005625
Exp. background	1.082161	0.062411	1.637963	0.087585	1.710763	0.103629
Observed	1	0	0	0	5	0

Table 8.4: Results of the counting experiment with the 1/3 data set, with all systematic parameters fixed to their mean values.

Phase	Run/GTID	T (MeV)	β_{14}	ITR	(x, y, z) (cm)	r (cm)	r_{NCD} (cm)
I	14438/0xc95cc	14.35	0.677	0.594	(-61.7, 13.0, 504.7)	508.6	—
III	51406/0x11bbbb	14.43	0.479	0.730	(-142.8, -178.9, -26.4)	230.4	34.0
III	54357/0x5a8af	15.35	0.527	0.795	(-69.9, 361.9, 211.1)	424.8	30.2
III	57484/0x155b96	15.58	0.482	0.810	(-254.9, -187.8, 188.3)	368.3	44.0
III	65318/0xbd463	15.14	0.666	0.791	(183.7, -194.9, 253.7)	430.1	62.5
III	65914/0xcc504	16.33	0.576	0.754	(42.9, 475.5, -176.5)	509.0	137.1

Table 8.5: Characteristics of observed events in the 1/3 data set. r_{NCD} is the distance to the closest NCD.

8.5 1/3 Dataset Results

This section summarizes the results of the *hep* and DSNB counting experiment analysis, accounting for the variation in systematic parameters. The joint distributions obtained by sampling systematics are shown in Figures 8.18 (Phase I), 8.19 (Phase II), and 8.20 (Phase III), and the signal and background expectations given in Table 8.6. In order to evaluate the relative impact of each systematic parameter on the total uncertainty, the standard deviation in signal and background expectation was calculated by varying each individually, with others held fixed at their respective means. These are given in Table 8.7; note that the total is different from the quadrature sum due to correlations. The dominant systematics for the *hep* and DSNB searches are very different: for the former, energy response is the major issue due to the steeply-falling ${}^8\text{B}$ solar neutrino background, and for the latter, the large flux and cross section uncertainties for atmospheric neutrinos dominate. These searches, however, remain very much statistically limited.

In total six events are observed in the *hep* ROI in the 1/3 data, one in Phase I and five in Phase III; the properties of these events are consistent with single electrons, and given in Table 8.5. These events have been hand-scanned to look for any unusual characteristics, and appear in all respects to be valid single-electron events. A partitioning of events so nonuniform in phase live time is unexpected, but consistent with statistic fluctuations in the small number of events involved. For example, the probability of observing ≥ 5 events in any one phase, ≤ 1 in a second phase, and zero in the other phase is around 7%. The fact that the combined three-phase background expectation is entirely consistent with the total observation, while the same detector (in a similar configuration) and the same underlying Monte Carlo model is used in all three phases, lends some credence to the interpretation that this is merely a somewhat unlikely fluctuation.

The least likely facet of this observation is the apparent upward fluctuation in Phase III, which we shall scrutinize further. A potential concern for that detector configuration is shadowing due to the NCDs: if an event mis-reconstructs near a detector, where the

ratio of energy to hits is expected to be smaller, the event will be assigned too high an energy. Figure 8.17 shows the reconstructed positions and directions of the events observed in Phase III (projected onto the $x - y$ plane) relative to the NCD array; there is no significant clustering around NCDs (or in z). There may be a spatial correlation with the outside edge of the NCD array, however this is difficult to establish with such a small sample. It seems unlikely that there is a global mis-modeling of the energy scale, resolution, or nonlinearity, since those parameters are measured using calibration data and any effect would have to be very dramatic to account for the observation. Non-Gaussian tails in the shape of the energy response function (due for example to the changes in detector optics introduced with the NCD array) are disfavored by the ${}^8\text{Li}$ studies in Section 7.4. Although that study relies on source data which is not available for positions very close to NCDs, data at $x = \pm 385$ cm is near the outer edge of the NCD array and shows no anomalies in energy reconstruction. Furthermore, there is no significant variation in 2D radial ($r = \sqrt{x^2 + y^2}$) or energy bias as a function of r in Monte Carlo; that is, if we bin the detector by concentric hollow diamonds following the shape of the NCD array, there is no bias introduced as we near the NCD array edge. This evidence supports the hypothesis that the ‘excess’ in Phase III is a normal fluctuation. In any case, it is important to bear in mind that somehow excluding these events worsens the agreement of the combined three-phase dataset with expectations.

An upper limit is computed for each three-phase pseudo-experiment in the ensemble (sampled with correlated systematic parameters), and the median is taken as the limit for the 1/3 dataset. The total signal and background expectations in the *hep* ROI are 1.077 ± 0.050 and 4.455 ± 0.429 , respectively, with six events observed. The sensitivity (the median upper limit for a signal-free model) is 4.19 times the BS05(OP) SSM prediction for the *hep* flux, or $3.32 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}$. For six events observed, an upper limit of 6.51 times the SSM is obtained, or $5.17 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}$. For comparison, the corresponding sensitivity for the previous Phase I analysis is 4.34 times the SSM, based on 0.99 ± 0.09 signal and

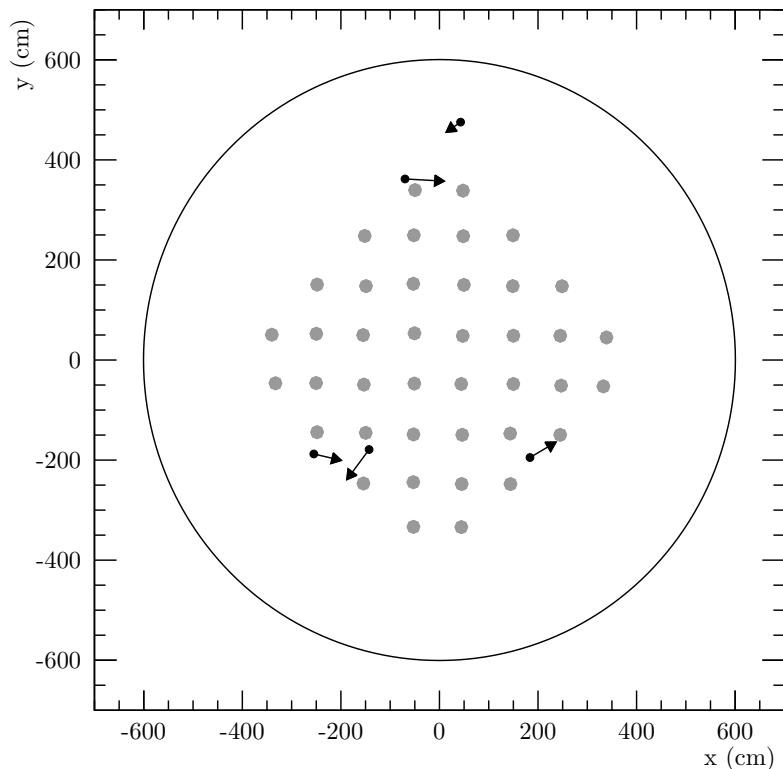


Figure 8.17: Reconstructed positions and directions of events observed in the *hep* ROI in the 1/3 dataset ($x - y$ projection). NCDs, shown as gray circles, are enlarged by a factor of 5. The arrows (which are arbitrarily but uniformly scaled) represent the reconstructed event direction, projected onto the $x - y$ plane. Further details about these events are listed in Table 8.5.

3.13 ± 0.60 background events expected; an upper limit is set at 2.9 times the BP2000 SSM (2.47 times BS05(OP), or $2.3 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}$) based on two events observed [15]. Note that the strong limit in that case is due to an apparent downward fluctuation, with fewer events observed than predicted for the background estimate alone, such that the limit is significantly stronger than the sensitivity. There is also some evidence that the high-levels cuts chosen for that analysis, in particular $\text{ITR} > 0.65$, could result in a biased estimate of the background rejection efficiency (see Section 8.1). For these reasons, we should not expect the limit to be more stringent than the previous value, even though the sensitivity is slightly improved.

In the DSNB ROI, no events are observed, with 0.03 signal and 0.70 ± 0.17 background events expected. With a background expectation close to unity, the distribution of Feldman-Cousins upper limits for an ensemble of Poisson-distributed pseudo-experiments is bimodal, depending on whether zero or one events are observed. Here, a median upper limit of 110 times the $T = 6$ MeV Beacom & Strigari model prediction lies between these extremes, which correspond to about 60 times the model for zero events, and 115 times the model for one event. With no events observed in the data I set an upper limit of 61.6 times the model prediction, which is fully consistent with the sensitivity. The model predicts $0.66 \nu \text{ cm}^{-2} \text{ s}^{-1}$ in the range $22.9 < T < 36.9$ MeV, and so the limit from the 1/3 data set corresponds to $\Phi_{\nu_e}^{\text{DSNB}} < 41 \text{ cm}^{-2} \text{ s}^{-1}$ in that energy range, an improvement over the existing direct ν_e limit [15] with only 1/3 of the full dataset analyzed. This results from increased statistics and a larger energy window, but also a random fluctuation toward zero events when observing one would be almost as probable.

8.5.1 Projections for the Full Dataset

Moving from the non-blinded dataset to the full data, the statistics for the search will be tripled. For the *hep* search, the corresponding sensitivity is 2.19 times the BS(05)OP SSM

8.5 1/3 Dataset Results

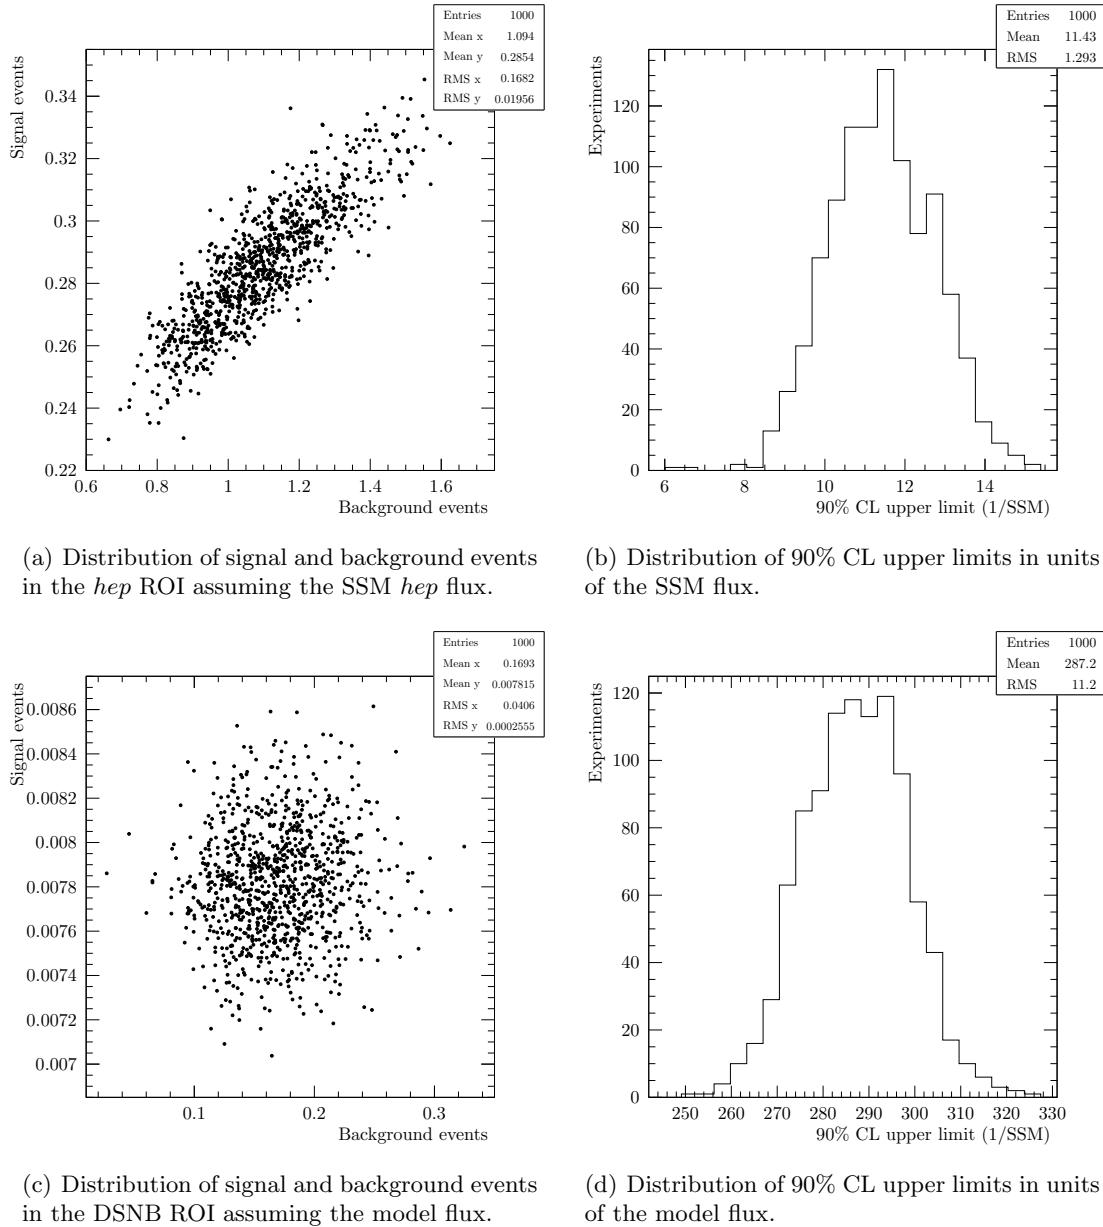


Figure 8.18: Monte Carlo systematic parameter variation in Phase I.

8.5 1/3 Dataset Results

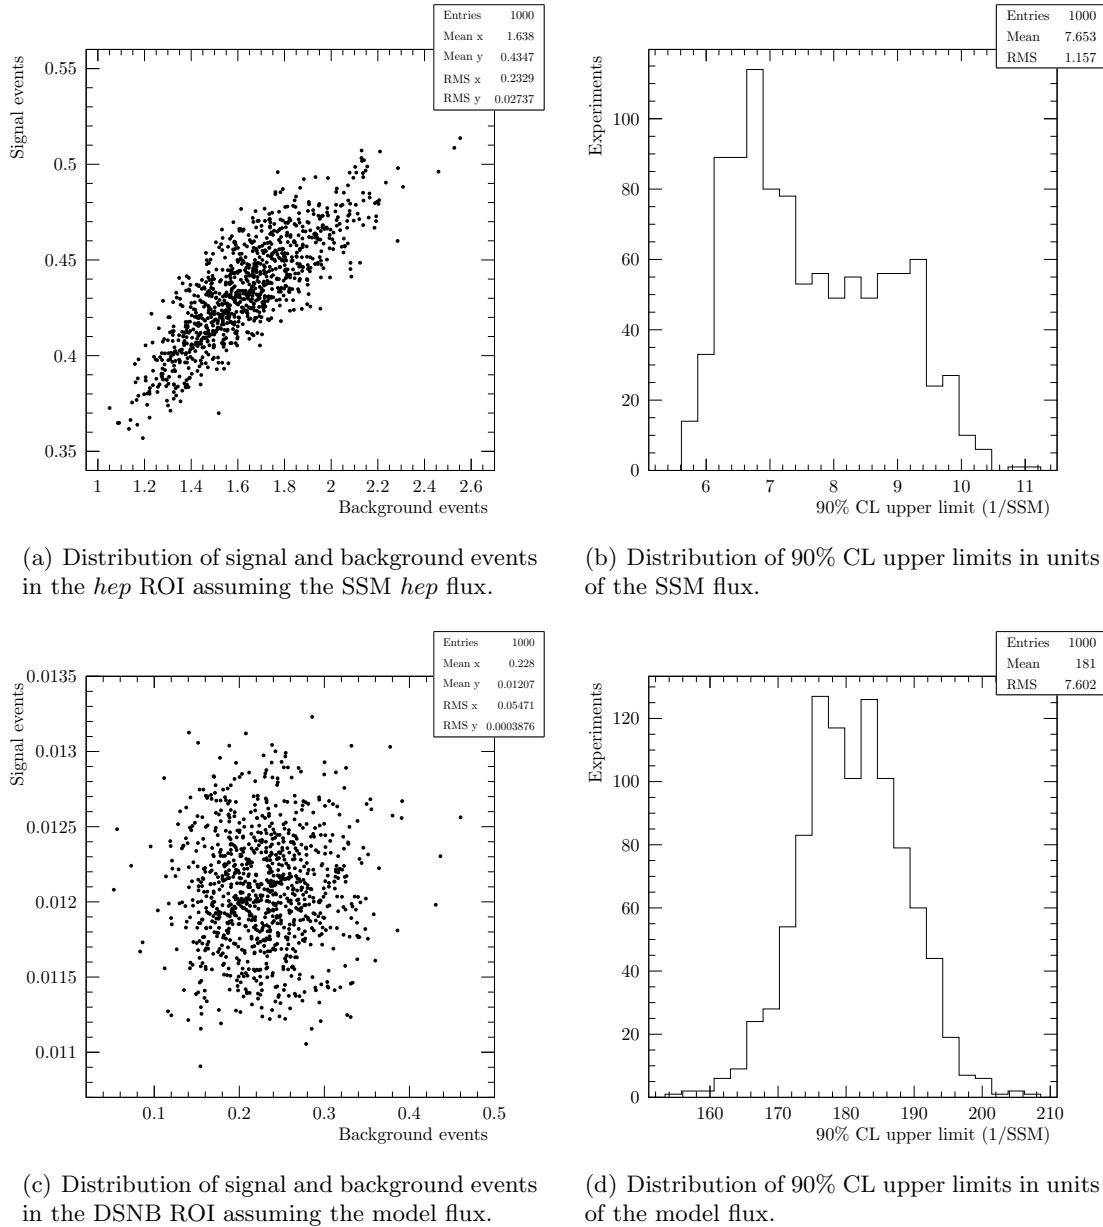


Figure 8.19: Monte Carlo systematic parameter variation in Phase II.

8.5 1/3 Dataset Results

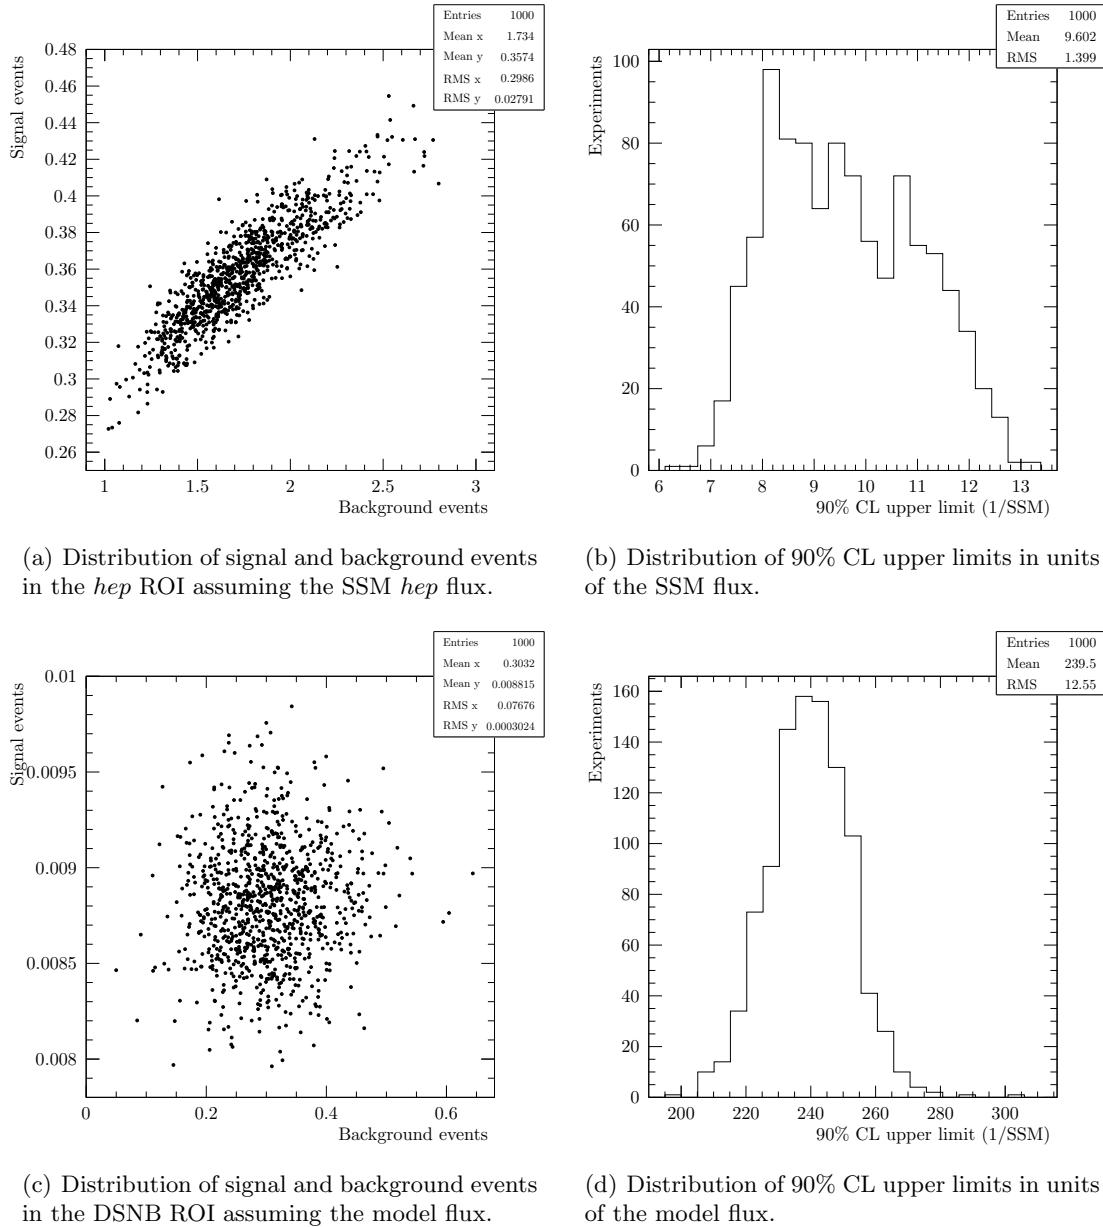


Figure 8.20: Monte Carlo systematic parameter variation in Phase III.

	Expected Signal	Expected Background	Events Observed
Phase I <i>hep</i>	0.283 ± 0.019	1.073 ± 0.161	1
Phase II <i>hep</i>	0.435 ± 0.027	1.639 ± 0.230	0
Phase III <i>hep</i>	0.359 ± 0.028	1.742 ± 0.287	5
Total <i>hep</i>	1.077 ± 0.050	4.455 ± 0.429	6
Phase I DSNB	0.008 ± 0.000	0.169 ± 0.041	0
Phase II DSNB	0.012 ± 0.000	0.228 ± 0.055	0
Phase III DSNB	0.009 ± 0.000	0.303 ± 0.077	0
Total DSNB	0.029 ± 0.001	0.700 ± 0.170	0

Table 8.6: Summary of systematics sampling results for the counting analysis signal and background expectations for the 1/3 data set.

prediction, or $1.74 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}$; this is the highest sensitivity to the *hep* flux of any search, and therefore we expect to obtain the most stringent limit. Given the proximity of the sensitivity to the model prediction, a low-significance first observation of this flux may also be possible, especially if the true value is near the upper end of the 15.5% model uncertainty.

The sensitivity for the DSNB flux with the full statistics is 42 times the Beacom & Strigari $T = 6$ MeV model prediction, corresponding to a ν_e flux of $28 \text{ cm}^{-2} \text{ s}^{-1}$ in the energy range $22.9 < T < 36.9$ MeV. This is significantly weaker than Super-Kamiokande’s limit on the $\bar{\nu}_e$ component of the DSNB flux ($1.4 - 1.9 \text{ cm}^{-2} \text{ s}^{-1}$) [14], but is by far the strongest direct limit on the ν_e component and improves upon even the model-dependent indirect limits derived from the SK data [71]. While a DSNB measurement is unlikely in light of the model predictions and the null result of SK, the improved limit can be used to constrain nonstandard models of supernova neutrino production and propagation. On the other hand, if some nonstandard model that enhances the ν_e flux is realized in nature, SNO may have the sensitivity to discover it!

	hep						DSNB					
	Phase I		Phase II		Phase III		Phase I		Phase II		Phase III	
	ΔS	ΔB										
Low-energy fit errors	0.00	4.90	0.00	5.39	0.00	5.77	0.00	0.00	0.00	0.00	0.00	0.02
Instrumental cut sacrifice	0.04	0.03	0.04	0.03	0.04	0.03	0.20	0.23	0.20	0.23	0.20	0.23
Live time	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Energy scale	5.44	11.65	3.98	9.02	6.82	14.62	0.78	4.01	0.51	4.09	1.11	3.23
Energy resolution	1.00	6.20	1.06	6.57	0.82	4.77	0.17	0.05	0.14	0.05	0.18	0.20
Vertex accuracy	2.72	2.72	2.78	2.78	2.51	2.51	2.72	2.72	2.78	2.78	2.52	2.51
Vertex resolution	0.04	0.08	0.04	0.10	0.04	0.06	0.05	0.58	0.04	0.78	0.05	0.46
${}^8\text{B}$ ν_e spectrum	0.00	0.01	0.00	0.01	0.00	0.01	0.00	0.00	0.00	0.00	0.00	0.00
θ_{12}	2.77	3.85	2.76	4.00	2.75	3.80	0.00	0.01	0.00	0.02	0.00	0.03
θ_{13}	0.23	0.22	0.23	0.23	0.23	0.22	0.00	0.00	0.00	0.00	0.00	0.00
Δm_{21}^2	0.43	0.19	0.43	0.21	0.44	0.20	0.00	0.00	0.00	0.00	0.00	0.00
Δm_{31}^2	0.01	0.01	0.01	0.01	0.01	0.01	0.00	0.00	0.00	0.00	0.00	0.00
Atm. ν flux												
$E_\nu > 100$ MeV	0.00	0.49	0.00	0.25	0.00	0.56	0.00	9.42	0.00	9.48	0.00	9.31
$E_\nu < 100$ MeV	0.00	0.02	0.00	0.01	0.00	0.02	0.00	0.90	0.00	0.71	0.00	1.05
Cross sections												
CC $\nu - d$	1.12	0.93	1.12	0.97	1.12	0.94	1.17	0.00	1.17	0.01	1.17	0.01
Atm. ν CCQE	0.00	0.25	0.00	0.19	0.00	0.28	0.00	12.62	0.00	11.74	0.00	8.11
Atm. ν other	0.00	1.02	0.00	0.50	0.00	1.05	0.00	14.39	0.00	15.34	0.00	19.27
15.1 MeV γ	0.00	0.52	0.00	0.15	0.00	1.05	0.00	0.00	0.00	0.00	0.00	0.00
Atm. ν n multiplicity	0.00	0.53	0.00	0.23	0.00	0.61	0.00	6.71	0.00	6.23	0.00	8.25

Table 8.7: Summary of the 1σ uncertainties on the signal (ΔS) and background (ΔB) expectations in the hep and DSNB ROI in each phase due to the variation of a single systematic parameter while others are fixed to their mean value. All units are percentages ($\Delta A = 100.0 \times \sigma_A/\mu_A$).

Chapter 9

Spectral Signal Extraction Fit

The counting analysis described in the previous chapter provides a straightforward approach to obtaining *hep* and DSNB limits from the SNO dataset, but does not take full advantage of the differences in the detected energy spectra of *hep* and ${}^8\text{B}$ solar neutrinos. This chapter describes a fit to the data, varying the normalizations and shapes of Monte Carlo probability distributions to extract the rate of *hep* neutrinos in the dataset. This approach improves the measurement of systematics and increases the statistics beyond the previously-defined energy region of interest as, in a sense, the fit naturally optimizes the ROI.

9.1 Signal Extraction

The goal of signal extraction is to determine the unknown parameters of a model given an observation of data. The signals and backgrounds are distributed in observable space (energy, position, particle ID, etc.), and corresponding probability distributions may be

9.2 The Maximum Likelihood Method

built using, for example, ancillary experiments or Monte Carlo techniques. In the case of SNO, a Monte Carlo model is validated and corrected based on comparison to calibration data. These probability distributions constitute a model which is combined with observed data to measure signal rates and detector parameters.

Drawing statistical inferences when fitting a model to data requires care, particularly when the model is complex (e.g. has a large number of highly correlated signals) or there are physical boundaries on the parameters one wishes to respect. These issues tend to drive the problem out of the Gaussian regime, invalidating many common shortcuts taken in function minimization.

Signal extraction for the present analysis faces both of these problems. The similar shapes of the signal and backgrounds result in correlations, and the measured signal rate is expected to be very close to the physical limit of zero. Hence, those function minimization techniques which assume smooth Gaussian-distributed parameters must be avoided when estimating model parameters. To this end, a maximum-likelihood method (Section 9.2) is used, with the maximization performed with a Markov Chain Monte Carlo sampling approach (Section 9.3), and parameter uncertainties extracted using both Bayesian and frequentist approaches (Section 9.4).

9.2 The Maximum Likelihood Method

The maximum likelihood (ML) method provides a means to determine the parameters in the model of the experiment; it is a tool for estimating the set of parameters that make the observed dataset the most likely to have occurred. One postulates a set of M signals, i.e. probability distribution functions from which the events were drawn, and the parameters describing those PDFs (e.g. normalization, energy scale) become ML *estimators* determined by maximizing a likelihood function. A dataset consists of N events, each described by a vector of observables \mathbf{x} .

9.2 The Maximum Likelihood Method

The sum \tilde{N} of the expectation values \tilde{N}_j for each signal yields the expectation value for the total number of events, around which the observed number N is Poisson-distributed. The probability for an individual event is then given by

$$P(\mathbf{x}|\{\tilde{N}_j\}) = \prod_{j=1}^M \frac{\tilde{N}_j}{\tilde{N}} \times P_j(\mathbf{x}). \quad (9.1)$$

Constructing a product over events to build a likelihood function, taking the logarithm, and dropping constant terms, one arrives at the negative log likelihood function, or NLL:

$$-\log \mathcal{L}(\{\tilde{N}_j\}) = \sum_{j=1}^M \tilde{N}_j - \sum_{i=1}^N \log \left(\sum_{j=1}^M \tilde{N}_j \times P_j(\mathbf{x}_i) \right). \quad (9.2)$$

Minimizing this function provides an estimate of the normalizations of the various signals that contribute to the dataset. Extending this method to include parameter constraints (adding penalty terms to Equation 9.2) or floating systematic uncertainties (letting P depend on additional parameters) is straightforward, making the ML approach quite flexible. The NLL may also encode prior knowledge about the parameters, such as the fact the event rates cannot be negative, if such constraints are desired. Furthermore, one is in general free to re-parameterize the NLL in terms of normalizations that mix signal rates, which may be desirable for including correlations between data sets or reducing the number of parameters. Including these terms in the NLL, one obtains the complete functional form suitable for signal extraction applications [27]:

$$\begin{aligned}
 -\log \mathcal{L}(\mathbf{r}, \Delta) = & \\
 & \sum_{j=1}^M \tilde{N}_j(\mathbf{r}, \Delta) \\
 & - \sum_{i=1}^N \log \left(\sum_{j=1}^M \tilde{N}_j(\mathbf{r}, \Delta) \times P_j(\mathbf{x}_i, \Delta) \right) \\
 & + \frac{1}{2} \sum_{k=1}^{M'} \frac{(r_k - \bar{r}_k)^2}{\sigma_{r_k}^2} \\
 & + \frac{1}{2} \sum_{m=1}^s \frac{(\Delta_m - \bar{\Delta}_m)^2}{\sigma_{\Delta_m}^2}
 \end{aligned} \tag{9.3}$$

where Δ represents the set of s systematic parameters with Gaussian uncertainties σ_Δ , and \mathbf{r} represents the rate parameters (related to \tilde{N} by a transformation matrix such that $\tilde{N}_i = \epsilon_i^j r_j$) and having Gaussian uncertainties σ_r . M' is simply the number of rate parameters which are externally constrained. Extension of the unbinned ML method to arbitrary observable dimensions is trivial, as a minimization algorithm sees only the probability $P_j(\mathbf{x}_i)$ and is unaware of the dimensionality of \mathbf{x} .

The likelihood space, with dimension equal to the number of parameters ($M + s$), contains complete knowledge of parameters given the observed data, including correlations to all orders. Given this function, one may calculate extrema as well as uncertainties on parameters, however defined. Simply calculating the entire space (e.g. by grid scanning) is impractical in high dimensions, so directed-search algorithms such as gradient descent are commonly employed to find extrema. These algorithms improve performance by making assumptions about the continuity and shape of the likelihood space, assumptions which break down near physical boundaries. Unfortunately, this is precisely the regime of interest in sensitivity studies and rare-process signal extraction. Markov Chain Monte Carlo

(MCMC), introduced in Section 9.3, provides a compromise between directed searches and grid sampling, trading some performance for robustness.

9.2.1 Floating Systematics

When we allow systematic parameters Δ to vary (float) in the fit, this will in general affect the number of events included in the fit range, as PDF shape distortions move events across the boundaries on observables which are fixed *a priori*. Naturally, fewer events in the analysis window V due to a systematic shift does not imply, e.g., a lower flux, and so \tilde{N} must be reweighted according to this efficiency. The Monte Carlo used to build the PDFs for this analysis is generated with a fixed scaling η relative to the model prediction (for example, 1000 times the BP2000 SSM) and the fit is in terms of a scaling in the total number of Monte Carlo events on top of that, such that I fit directly for a scaling relative to the prediction of the model (which is easily converted into a flux). There is therefore no need to normalize based on the number of events in V without any transformation applied; this is special case which must cancel out anyway, as the answer cannot depend on some arbitrary (if trivial) choice of systematic parameters. The elements for signal i in the transformation matrix ϵ include this weighting with a factor of $|\{\mathbf{x}_i | S(\mathbf{x}_i, \Delta) \in V\}|$, which counts the number of events in the analysis window (a volume V in the observable space) after application of a systematic transformation S .

9.2.2 Multi-Phase Fitting

I perform a simultaneous fit to the entire three-phase data set, rather than combining the results of separate fits, in order to better constrain the ${}^8\text{B}$ and *hep* fluxes (and to a lesser extent the atmospheric neutrino background): since these sources are presumably constant across all three phases, the expected rates should scale in the same way. Using the combined statistics of all three phases helps to reduce the uncertainty in this background, which in

the counting analysis was second only to energy-related systematics in driving the total systematic uncertainty (see Table 8.7).

Phase information is included in the signal extraction fit by adding a “dataset” tag to each event vector \mathbf{x}_i , and only evaluating the (phase-specific) PDFs for the appropriate events. Correlations of rates are handled as noted above, by replacing \tilde{N}_j with a product of a source rate (which is the fit parameter) and a normalization matrix (which accounts for the differences between phases in the expected numbers of events for a given flux).

9.3 Markov Chain Monte Carlo

Markov Chain Monte Carlo provides a general technique for sampling probability distributions, with broad applications. Essentially, the domain is sampled by a random walk such that the distribution of random steps approximates the probability distribution itself. The Markov Chain refers to a sequence of steps in a discrete-time stochastic Markov process, i.e. generated by transitions within a set of states obeying the Markov property (see, e.g. Reference [111]). This defining property is that each transition depends only on the current state (not previous history), however the following conditions are also desirable for tractable systems:

Stationary The system has a steady-state solution

Unique There exists one and only one stationary solution

Ergodic The system is aperiodic (no “infinite loops” in the graph) and positive-recurrent (the graph is completely connected)

For a discrete state space, we can envision a graph like in Figure 9.1, with transition

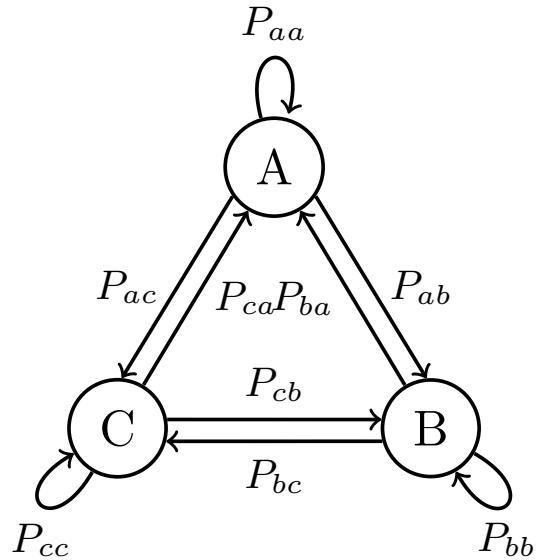


Figure 9.1: A three-state system.

probabilities representable in a matrix form:

$$Q = \begin{bmatrix} P_{aa} & P_{ab} & P_{ac} \\ P_{ba} & P_{bb} & P_{bc} \\ P_{ca} & P_{cb} & P_{cc} \end{bmatrix}.$$

For a state vector x^i , the probabilities for the next state are given by $x^{i+1} = Qx^i$, and if the above conditions are satisfied, an equilibrium (stationary state) is reached after many steps:

$$\lim_{N \rightarrow \infty} Q^N = \begin{bmatrix} P_a & P_b & P_c \\ P_a & P_b & P_c \\ P_a & P_b & P_c \end{bmatrix}$$

where columns give the probability to be in each state, independent of the initial conditions. This example applies to problems where the state space is discrete and the transition

probabilities are analytically calculable. In our case, where we seek to map out an unknown probability distribution with continuous support, we choose initial conditions and assert the process dynamics, and perform many transitions in a numerical simulation to discover the steady state.

The applications of the Markov Chain Monte Carlo technique to sampling arbitrary probability distributions were elucidated by Metropolis et al. in a landmark 1953 paper [112], where the authors developed the transition algorithm in analogy with statistical mechanics. The Metropolis algorithm is as follows:

1. Propose a step to coordinate \vec{x} from current coordinate \vec{x}_0 based on a symmetric sampling distribution $Q(\vec{x}|\vec{x}_0)$.
2. If the probability $P(\vec{x}) > P(\vec{x}_0)$, or if the ratio of probabilities $P(\vec{x})/P(\vec{x}_0) \geq X$, where X is uniformly distributed on $[0, 1)$, move to \vec{x} . Otherwise, remain at \vec{x}_0 .
3. Record the current coordinate.

Many extensions to the Metropolis algorithm have been developed, most notably Metropolis-Hastings [113], which allows for asymmetric proposal distributions by including a weighting factor. Beyond this, almost all extensions are special cases of Metropolis-Hastings, applicable to a subset of problems. For generality, ease of implementation, and because there is no particular motivation to use an asymmetric proposal distribution, I use the basic Metropolis algorithm in the subsequent analysis.

In the limit of an infinite number of steps, coordinates will be visited with a frequency according to the underlying probability distribution P , and so, given the list of sampled points, one is able to extract the properties of the distribution. This approach is particularly powerful because it is suitable for spaces of any dimension, and trivially handles complicated or non-Gaussian distributions. By using an MCMC to sample a likelihood function, we may

leverage the flexibility of this method to build a robust parameter estimator. Crucially, however, the sampled distribution is not the same as $P(\vec{x})$, but is only proportional to $P(\vec{x})$. One may normalize the distribution by choosing a prior distribution and applying Bayes' Theorem (see e.g. [108]), to obtain a posterior probability distribution from which Bayesian credible intervals may be derived. However, it is not possible to apply the (frequentist) profile likelihood approach here without an absolute measure of the likelihood. For this reason the above procedure is augmented to also record the value of $P(\vec{x})$ for all sampled states $\{\vec{x}\}$. In this way, the MCMC becomes a form of importance-weighted sampling, providing a non-uniformly sampled version of P with a sample density proportional to the probability. The sampled set $\{P(\vec{x})\}$ may be used to calculate frequentist intervals.

There are two additional complications to the MCMC approach. First, we must choose a starting point according to the (as yet unknown) distribution. Consider a case where in the above discrete-state example, we chose an initial state A which actually had a vanishingly small probability. For a finite set of samples, the value obtained for P_a would be too large. We avoid this situation through the use of a “burn-in” phase, discarding the initial steps so that the initial conditions are irrelevant. Second, there is the matter of ensuring that the chain is well-mixed and that the fit has converged to a steady state.

9.3.1 Convergence

In general, it is impossible to know with certainty that the MCMC has converged. For example one can easily imagine sampling only around a local minimum in a much larger parameter space, if such a minimum existed and the initial proposal distribution was chosen poorly (too narrow). In our case, however, it is known that the parameters are roughly Poisson- (normalizations) or Gaussian-distributed (shape systematics, as evidenced by fits to source data). Furthermore, the estimates of central values are very robust. Therefore the likelihood function too should approximately follow a normal distribution (truncated

in the case of parameters near a boundary), and when we arrive at a stationary posterior distribution, we can be confident that we have obtained a good estimate of the likelihood space.

A wide array of formal convergence tests have been developed, with a few notable examples due to Gelman and Rubin [114] and Raftery and Lewis [115]. For a review of commonly-used approaches, see Reference [116]. With no perfect method, however, the criteria remain subjective at some level, with the most robust tool being inspection of the parameter time series and the posterior distribution itself, as illustrated in Figure 9.2. To formalize this process, I break the samples for each parameter into N time slices, and compute an unbinned Kolmogorov-Smirnov (KS) test comparing each slice to each other slice; this returns a probability that the samples in question were drawn from the same parent distribution, which should be the case if we have reached the stationary distribution.

The autocorrelation function for the time series of each parameter also provides a measure of mixing: for an appropriately-chosen proposal distribution, samples should rapidly lose correlation as a function of the time separation (known as the lag).

9.4 Estimating Uncertainty

MCMC provides a robust technique for determining maximum-likelihood estimators for likelihood functions containing physical boundaries, excluded regions, or local maxima. However, this is only a part of the signal extraction puzzle: also of great importance are the uncertainties on the parameters.

The meaning of uncertainty is itself uncertain; a variety of methods which differ both algorithmically and philosophically are in common use in particle physics. Both objective confidence intervals calculated using a frequentist framework and subjective credible intervals derived following a Bayesian approach will be presented.

9.4 Estimating Uncertainty

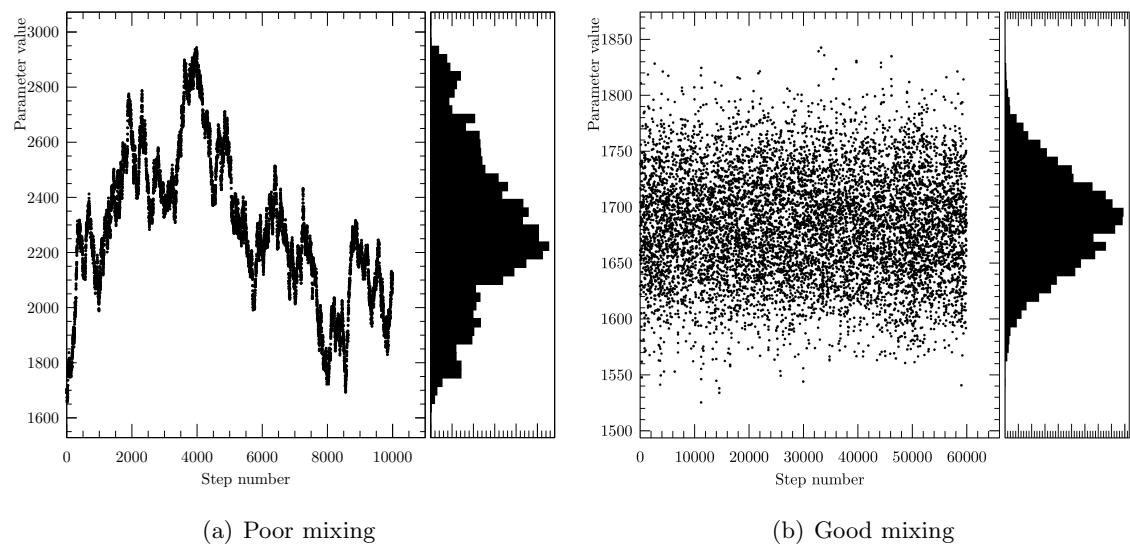


Figure 9.2: Examples of parameter time series and projected distributions for a poorly-mixed and well-mixed chain.

9.4.1 The Bayesian Interpretation

One approach to determining parameter errors given a sampled likelihood space is to treat the space like a posterior probability distribution for the parameters. To map the likelihood function to a probability distribution, one must normalize by assuming a prior distribution for the parameters. The choice of Bayesian prior is not obvious for a Poisson process; see Reference [117] for a discussion. Here, I choose a uniform (“flat”) prior (the canonical “uninformative” choice) for physically-allowed regions of the parameter space (i.e. Poisson means are ≥ 0).

When treating the normalized likelihood space as a probability distribution, one is free to integrate it directly to obtain credible intervals containing a desired fraction of the total probability. I project out each dimension from the full n -dimensional likelihood space and find intervals as close to central as possible; that is, for a confidence level $1 - \alpha$, parameters (θ, \mathbf{r}) where θ represents the parameters of interest and \mathbf{r} the other (nuisance) parameters, an interval $[\mu_1, \mu_2]$ is chosen such that

$$\frac{\int_{\mu_1}^{\mu_2} \int \cdots \int_{V_{\mathbf{r}}} \mathcal{L} d\theta d\mathbf{r}}{\int \cdots \int_{V_{(\theta, \mathbf{r})}} \mathcal{L} d\theta d\mathbf{r}} = 1 - \alpha, \quad (9.4)$$

and

$${}_{\theta} \int_{\mu_1}^{\hat{\mu}} P_{\theta} d\theta = {}_{\theta} \int_{\hat{\mu}}^{\mu_2} P_{\theta} d\theta = (1 - \alpha)/2 \quad (9.5)$$

where $\hat{\mu}$ is the mean of a Gaussian fit to the normalized marginal distribution P_{θ} for parameter θ . In cases where such a central interval does not exist (for instance, for parameter values near a physical boundary), we find the one-sided interval $[\mu_1 = 0, \mu_2]$ such that

$$\int_0^{\mu_2} P d\theta = 1 - \alpha. \quad (9.6)$$

It is important to note that these credible intervals cannot be compared directly to frequentist confidence intervals, including those derived in the counting analysis; these objects have very different meanings. Numerically, one-sided credible intervals constructed in this way will have an upper limit $\Phi^{-1}(1 - (1 - \alpha)/2)/\Phi^{-1}(\alpha)$ smaller than the equivalent Feldman-Cousins intervals, where $\Phi^{-1}(\alpha)$ is the inverse quantile function, due to the built-in “flip-flopping” protection in the unified frequentist approach. In the case of a normally-distributed parameter this is the probit function and the decrease is 28.2%, which must not be mistaken for any sort of improvement in the limit. Also, in the Bayesian construction there is no guarantee of coverage: 90% credible intervals will not necessarily contain the true value for 90% of experiments measuring that value. We may demonstrate explicitly that Bayesian limits constructed in this way do have proper coverage in the parameter range of interest using a Monte Carlo approach, though must concede that true Bayesians would deem this fact irrelevant.

9.4.2 The Profile Likelihood Construction

According to the frequentist interpretation, the likelihood space cannot be treated as a probability distribution: there exists a distribution of possible measurements given the true parameter values μ , $P(x|\mu)$, but the values of the parameters are fixed by Nature and have no probability distribution $P(\mu|x)$. Uncertainties therefore are defined in terms of an ensemble, such that the desired frequentist coverage is assured. In other words, the interval for a particular measurement is a randomly-distributed range in parameter space, which will only contain the true value in some fixed fraction of trials.

According to the Neyman construction [118], we use the multi-dimensional distributions $P(x, \mu)$ (here derived from Monte Carlo) at physically-allowed values of μ to define an acceptance region in the parameter space within which the coverage requirement is satisfied. In analogy with hypothesis testing, the acceptance region is defined by a likelihood ratio

9.4 Estimating Uncertainty

ordering principle, as defined in Appendix G:

$$\tilde{R}(\mu_s) = \frac{\mathcal{L}(x|\mu_s, \hat{\mu})}{\mathcal{L}(x|\hat{\mu}_s, \hat{\mu})} \quad (9.7)$$

where μ_s is a fixed value of the parameter of interest and $\hat{\mu}$ conditionally maximizes the likelihood with fixed μ_s . The set of parameters $\{\hat{\mu}_s, \hat{\mu}\}$ maximize the likelihood globally. This projects out the μ_s dimension of the likelihood space along the line of maximum likelihood, thereby avoiding integration over the distributions of “nuisance” parameters μ , as is done in the Bayesian case.

In principle, to calculate the sensitivity one must explicitly construct the acceptance region in the parameter space for each of an ensemble of fake datasets D assuming some fixed hypothesis, by using Monte Carlo simulation of a secondary ensemble of fake datasets $\{F\}$ sampled from the best-fit physically-allowed parameters of a fit to D . In practice, however, the procedure is considerably simpler in the limit of large N . If the likelihood function is normal or, by the property of invariance, there exists a transformation of variables f through which \mathcal{L} is made normal, an interval defined by R such that

$$\log \mathcal{L}(D|\theta, r) = \log \mathcal{L}_{\max} - \frac{1}{2}\chi^2(\beta, k) \quad (9.8)$$

in k dimensions will have probability content

$$P(-2 \log R \leq \chi^2(\beta, k)) = \beta$$

since $R = -2 \log(\mathcal{L}(D|\theta)/\mathcal{L}(D|\hat{\theta}))$ is asymptotically distributed as $\chi^2(\beta, k)$ [108]. To evaluate single-parameter uncertainties, for example, we use $\Delta(\log \mathcal{L}) = \chi^2(\beta = 0.683, dof = 1)/2 = 1/2$ (“1 σ ” errors) and $\chi^2(0.9, 1)/2 = 2.71/2$ (90% CL errors).

In the present analysis, parameters representing the Poisson mean of signal rates are defined to be non-negative. If such a parameter is very close to zero, the likelihood space is highly non-Gaussian and the parameter transformation f correspondingly nonlinear. In order to verify that the first-order χ^2 approximation is valid and that this method produces viable frequentist intervals for the *hep* rate, I perform many fits to fake data with known signal rates from 0.1 to 10 times the SSM prediction (much larger than the 5σ model uncertainty) and verify that the expected coverage is obtained in the ensemble limit in the range of interest for the *hep* signal rate. This ensemble testing is discussed further in Section 9.7.

9.4.3 A Two-dimensional Example

To illustrate the power of the MCMC method for signal extraction, let us consider an example with two signal PDFs, a flat background and a Gaussian signal. The likelihood function is defined so that rates are ≥ 0 via a severe penalty term. To estimate the sensitivity of this experiment, we perform an ensemble of fake experiments with data sets sampled from the background PDF. We then run the MCMC algorithm to map out the 2D likelihood space, and locate the maximum to determine the best-fit signal and background normalizations. Figure 9.3 shows an example fit.

The sensitivity is then computed in two ways: projecting the likelihood space onto the signal normalization dimension and finding the rate parameter value below which 90% of the likelihood falls (see Section 9.4.1), and using a profile likelihood approach (Section 9.4.2). The final sensitivity estimate $\langle \hat{S} \rangle$ is the median sensitivity from fits to an ensemble to independent fake data sets, with results shown in Figure 9.4. Note that although the numerical values of the sensitivities are similar, the meanings are entirely different; the profile method limits are effectively 95% limits due to the flip-flopping protection implicit in the likelihood ratio construction.

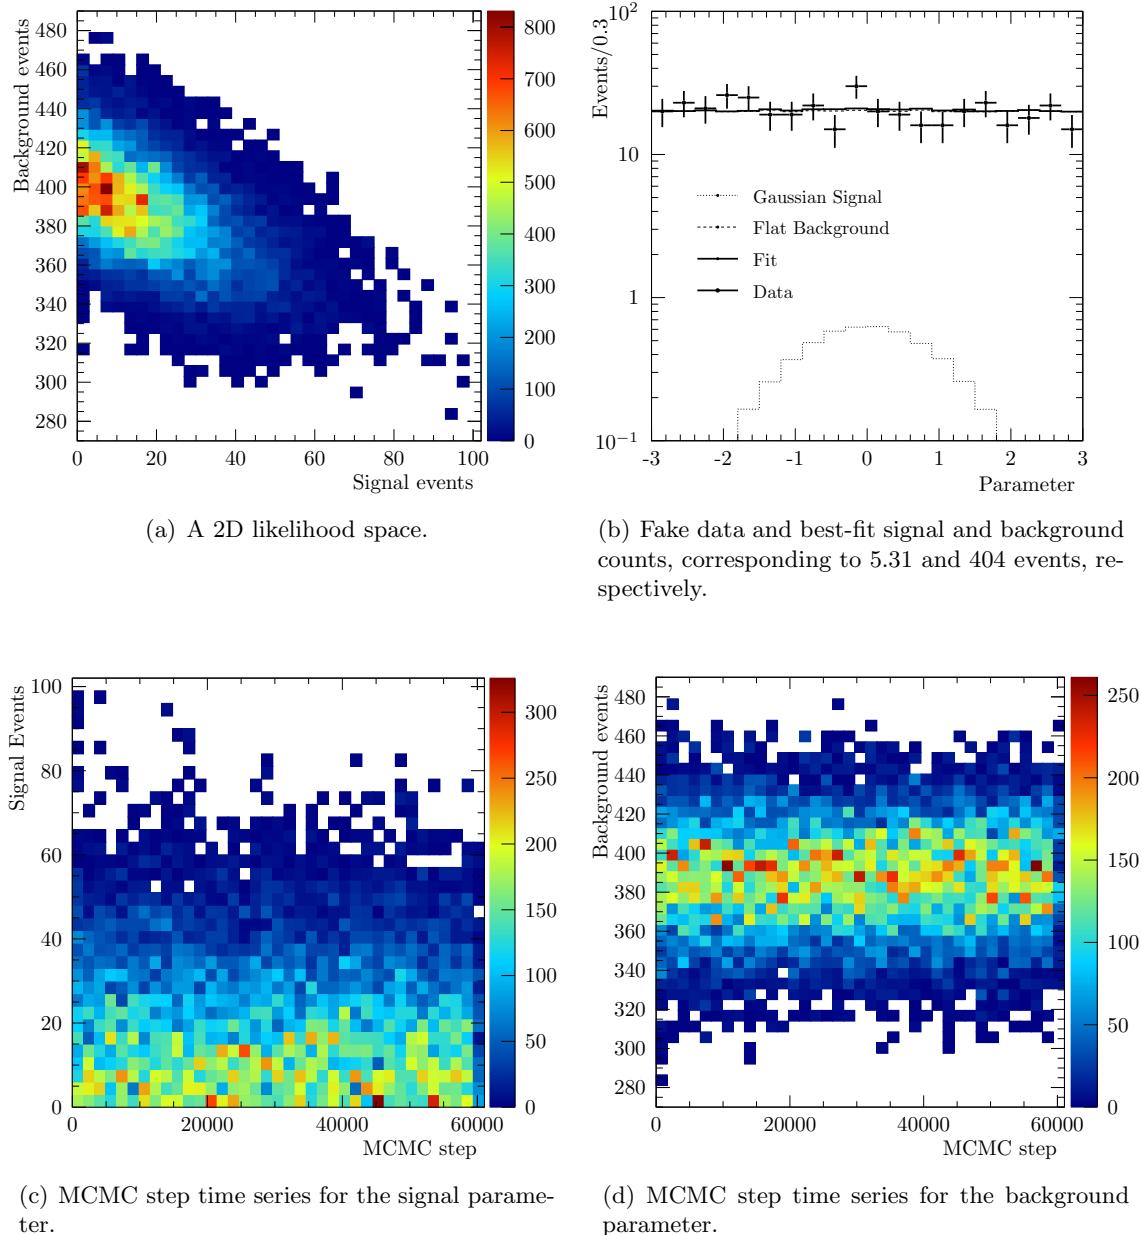


Figure 9.3: A fit to a signal-free fake experiment in a two-signal example.

9.4 Estimating Uncertainty

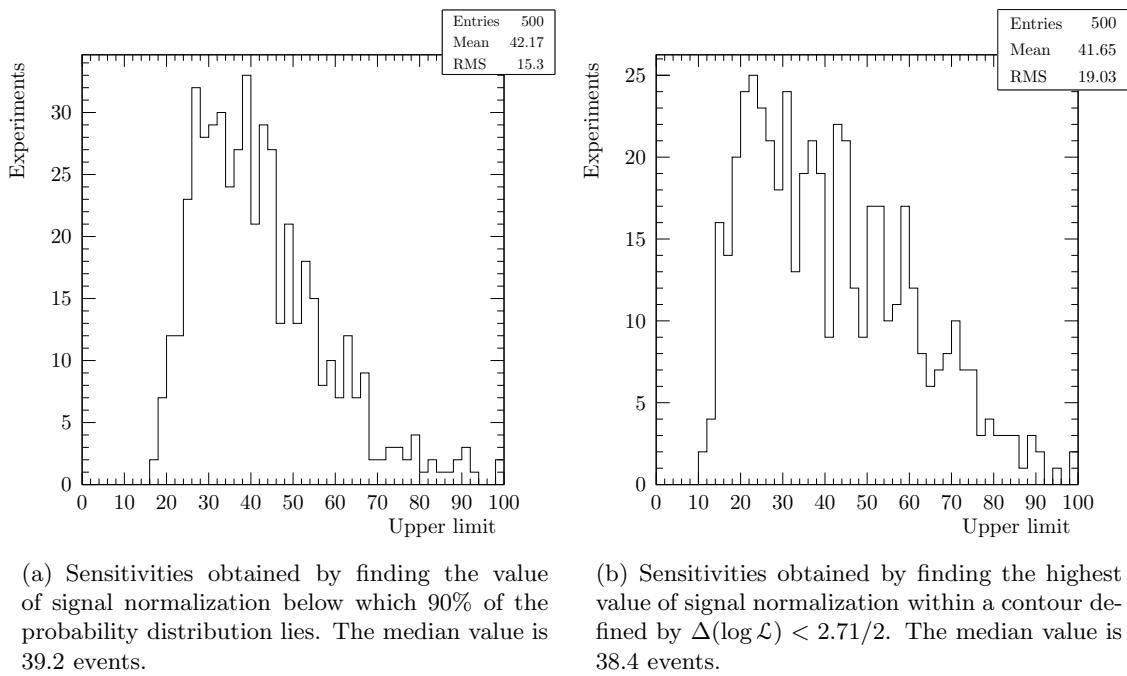


Figure 9.4: A comparison of sensitivities obtained by integration and the Feldman-Cousins approach.

Bias Near a Physical Boundary In order to verify the accuracy of the fit method, we typically would perform ensemble tests and generate bias and pull distributions to check for normality, centrality, and pull width. The one-sidedness of the likelihood function near a parameter boundary, however, means that the statistical fluctuations in the data or the location of the maximum always pull the signal mean upward. This introduces both variance and bias, and the pull distribution is not well-defined for parameters with asymmetric uncertainties. In the profile likelihood case, where coverage is a well-defined quantity, we may ensure that it is correct for ensembles of pseudo-experiments with a relevant range of signal rates. This form of ensemble testing ensures that the intervals returned by the fit are valid for the desired confidence level.

In the case of the *hep* flux in particular, we have a definite, nonzero model prediction for the expected signal. I therefore explicitly verify the coverage within e.g. the 5σ model uncertainties, and ensure that the fit produces valid intervals for any true value of the flux, by performing an ensemble of fits to Monte Carlo data sets.

9.5 The `sxmc` Code

To apply these MCMC-based ML fitting techniques to the SNO *hep* search, I have developed a software package called `sxmc`, short for Signal Extraction with a Markov Chain Monte Carlo [119]. This software builds PDF histograms from input Monte Carlo, and uses an MCMC to map out the likelihood function in arbitrary user-defined dimensions. `sxmc` allows the user to float n -th order scale, shift, and resolution scaling systematics for any observable. It includes a suite of statistical tools for making inferences including limit-setting, and tools for evaluating goodness-of-fit, Markov chain convergence, and bias. The major motivations for an MCMC-based fitter are that it is straightforward to apply parameter boundaries and that the algorithm makes only very modest assumptions about the properties of the underlying likelihood space, handling non-Gaussian functions and local minima well.

9.5.1 Implementation

`sxmc` is implemented in C++03 and Nvidia CUDA [120], and makes heavy use of general-purpose graphics processing units (GPGPUs) for hardware acceleration. GPUs, originally developed for computer graphics applications, provide an architecture highly optimized for large numbers of small parallel computations. These devices also have large on-board memory with excellent memory bandwidth, making them a powerful tool for parallel processing of large datasets. In the case of `sxmc`, GPGPUs are used in two computationally-intensive steps: histogram building and likelihood evaluation.

At each step, the likelihood function must be evaluated for the current parameter vector. As described above, `sxmc` computes an unbinned likelihood, and so this step involves iterating through all the events in the dataset being fit. The first optimization is building a two-dimensional lookup table (LUT). Every signal PDF is evaluated for each event; this provides the $P_j(\mathbf{x}_i)$ in Equation 9.3, and in the case that PDF shapes are constant, this computation only needs to be performed once. Given the LUT and a set of normalization parameters (i.e. \tilde{N}_j), the next step is to compute the second term in Equation 9.3 by looping over data events. This is performed in parallel on the GPU, in a staged sum that both preserves precision and ensures that GPU threads access memory sequentially, improving performance. Finally, this sum is combined with the normalization and constraint terms to arrive at the NLL. At this point the NLL remains on the GPU device, avoiding the significant time of overhead associated with device-to-host memory transfers. The next parameter step is chosen in GPU code and the process repeats, with parameter vectors and likelihood values being buffered in GPU memory. This buffer is periodically synchronized to host memory in efficient bulk transfers and flushed to a ROOT file for storage.

When systematic parameters are allowed to vary in the fit, the PDF histograms must be rebuilt from Monte Carlo for each step in the Markov Chain, taking into account the new systematic shift. In order to obtain smooth PDFs for the fit, the MC generally consists of million events, and so looping through them is quite time-consuming. This is mitigated

Platform	Samples/s	Notes
Intel Core i7 920	1.85×10^7	CPU mode, 2.67 GHz
Nvidia GeForce GT 650M	5.72×10^8	Apple MacBookPro10,1
Nvidia GeForce GTX 580	1.61×10^9	
Nvidia Tesla K40	3.00×10^9	

Table 9.1: Performance of `sxmc` PDF histogram building on a variety of platforms.

by moving the MC dataset to the device memory and building histograms in parallel on the GPU. A histogram class on the GPU keeps track of MC samples, descriptions of the systematic parameters, and performs on-demand re-binning of events when those parameters change.¹⁷. Table 9.1 shows the performance of PDF histogram building on a variety of architectures.

The CUDA GPGPU code is targeted to Nvidia hardware. However, `sxmc` is implemented using the `hemi` library [121], and so may also be compiled to run on traditional CPUs, with parallel operations serialized in loops (with a corresponding loss in performance).

9.5.2 Proposal Distributions and Burn-in

As noted above, two challenges that arise in implementing an MCMC are preventing the choice of initial conditions from biasing the final result, and achieving good mixing to efficiently explore the parameter space. In `sxmc`, these are both addressed in an initial two-part “burn-in” phase. Parameters are initially set to their expectation values, and width of the multivariate Gaussian sampling distribution $Q(\vec{x}'|\vec{x})$ are chosen in each dimension as:

$$\sigma_Q^i = \frac{\xi}{10} \sqrt{\max(\sigma^i, \tilde{\sigma}^i)}$$

¹⁷The `sxmc` GPU histogram code “pdfz” was originally developed by the legendary Stan Seibert.

or the constraint standard deviation if a constraint is provided, where $\tilde{\sigma}$ is 10 for normalizations and 1 for systematic parameters. The scale factor ξ is set to $(2.4)^2/N_{\text{params}}$ [122], and `sxmc` allows the user to optionally specify an additional scaling. After $\mathcal{O}(10^6)$ steps with this proposal distribution, the RMS of the accepted steps σ' is computed for each dimension, and the proposal width is set to $(\xi\sigma')$. This sequence is repeated once more to further refine the width, which is then fixed for the sampling used for the fit. This approach converges to a step acceptance rate of $\sim 20\%$ for a wide variety of fit conditions.

9.6 Signals and Backgrounds

The signal extraction fit includes 18 signals — ${}^8\text{B}$ CC, ${}^8\text{B}$ ES, *hep* CC, *hep* ES, neutrons, and atmospheric neutrino interactions for each of the three phases — through six free parameters: the ${}^8\text{B}$ flux scaling, *hep* flux scaling, atmospheric neutrino flux scaling, and independent neutron rates for each phase.

${}^8\text{B}$ and *hep* Electrons Electrons are produced in solar neutrino interactions in SNO via several channels: CC interactions on deuterium, ${}^{17}\text{O}$, and ${}^{18}\text{O}$, and ES on electrons in the target volume and the AV acrylic. The shape of the total spectrum depends systematically on the ratios of the cross sections for these properties, number of target nuclei or electrons, and the oscillation model which determines the $\nu_e/\nu_{\mu\tau}$ ratio at any given energy. Since the goal of this analysis is not to measure these ratios or mixing parameters, these quantities are fixed in the fit. PDFs are included in the fit for ${}^8\text{B}$ CC, ${}^8\text{B}$ ES, *hep* CC, and *hep* ES, but the relative normalizations of ${}^8\text{B}$ CC and ES signals and the *hep* CC and ES signals are fixed; the sole reason for the distinction is to apply different systematics to CC and ES. The uncertainty in the mixing parameters is also included as a systematic, as described below. The underlying ${}^8\text{B}$ rate and *hep* rate float freely in the fit, but are constrained to be equal across all three phases.

Neutrons The observed energy for neutrons is sufficiently low that it does not affect *hep* search directly, but is included as a nuisance parameter such that the fit may extend lower in energy, improving the constraint on the ${}^8\text{B}$ electron background. This signal absorbs all neutron captures in the detector, from NC interactions and other sources. Since this is not a direct solar neutrino NC measurement, this rate is considered to be independent in each phase.

Atmospheric Neutrino Interactions All atmospheric neutrino interaction channels are combined in the fit, and the summed spectrum is floated. There is uncertainty in the spectrum shape due to different cross section uncertainties for different channels; this is handled as described below. This rate is also assumed to be constant across the three phases, after scaling according to the time-averaged flux given the position in the solar cycle.

9.6.1 Cuts

For inclusion in the signal extraction fit, events must pass a number of cuts, as detailed in Section 8.1. The suite of cuts includes:

- Radius $r < 550$ cm
- ITR > 0.55
- $P_{2d} > 10^{-2}$
- $P_\phi > 10^{-6}$
- ICT $> 10^{-6}$

Energy and isotropy become observables in the signal extraction fit, in contrast to the counting analysis. The cuts are intended to minimize the atmospheric neutrino background in the fit range (see discussion in Section 8.1).

9.6.2 Observables

The observables included in the fit must be chosen carefully, to maximize the separation of signals while minimizing the dimensionality of the space so as not to dilute PDF statistics. For the *hep* fit, three observable dimensions are used: the energy, the angle of events relative to the Sun, and the isotropy parameter β_{14} . Rather than relying on an external constraint, this signal extraction fit effectively includes the entire low-energy signal extraction described in Section 8.2.

While it is evident from the HLC optimization that some of these parameters, ICT in particular, would potentially help discriminate signal and atmospherics in a fit, the actual statistics for the latter are too small to be of any use, and so there is little motivation to increase the dimensionality of the fit beyond what differentiates *hep* neutrinos, ${}^8\text{B}$ neutrinos, and neutrons.

Energy Energy is the main discriminant for separating *hep* from ${}^8\text{B}$ solar neutrino events. The quantity used in the fit is the reconstructed effective energy kinetic energy T_{eff} , which is the most likely kinetic energy under the hypothesis that the event was due to a single electron. As in the counting analysis, the best position fitter for each phase is used (the Path Fitter FTP for Phases I and II, and NCD-aware QPDF fitter nFTU in Phase III), and the energy response fitter RSP is used in all phases. These reconstruction algorithms are reviewed in Section 5.6.

T_{eff} is binned into fourteen 1 MeV bins ranging from 6–20 MeV, and so includes the entire counting experiment low-energy signal extraction range, the *hep* ROI, and the range in between. The T_{eff} spectra are shown in Figure 9.5(a).

Angle Relative to the Sun The reconstructed angle relative to the Sun ($\cos \theta_{\odot}$) separates neutron and atmospheric backgrounds from solar neutrinos since the former are uncorrelated with the solar direction while the (CC) ES component of the latter is (anti-)

9.6 Signals and Backgrounds

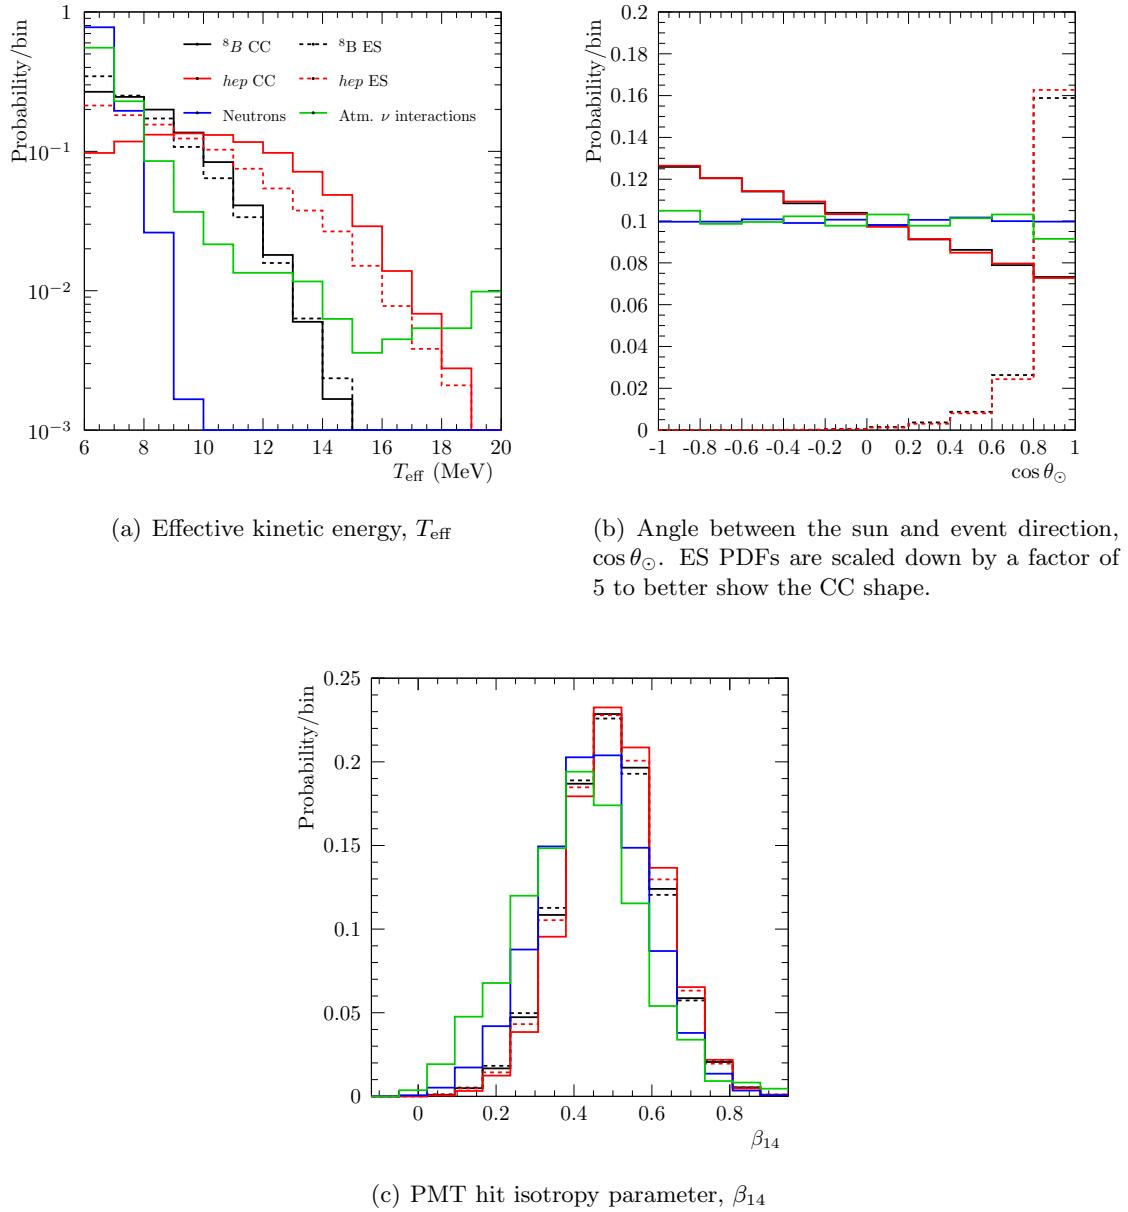


Figure 9.5: One-dimensional projections of the PDFs for Phase I.

correlated. This dimension uses 10 bins ranging from -1 to +1. The $\cos \theta_\odot$ spectra are shown in Figure 9.5(b). Interestingly, there is a small amount of discrimination between hep and 8B neutrinos afforded by the different fractions of CC and ES in the energy window, which distorts the total angular spectrum.

Isotropy β_{14} The β_{14} isotropy parameter helps to discriminate neutron from neutrino events, and also provides some handle on the atmospheric neutrino background. This observable is primarily included to better measure the neutron background to the 8B electrons at low energy. The fit uses 15 bins in the range $-0.12 < \beta_{14} < 0.95$. The β_{14} spectra are shown in Figure 9.5(c).

9.6.3 Systematic Uncertainties

The same systematic uncertainties are considered for the signal extraction fit as were for the counting analysis, and these are summarized in Table 6.4. For the fit, systematics are treated with one of two methods: floating and scanning.

In the case of floating systematics, the parameters are included in the fit and modify the PDF shapes; this simultaneously optimizes the signal shapes and normalizations. The parameters floated in the fits include an energy scaling and resolution scaling with the parameterization given in Section 7.2 and mean values and constraints given in Table 7.3. An angular resolution scaling (specifically as applied to $\cos \theta_\odot$) is applied to the ES signals, and parameterized following previous SNO results [46]:

$$(\cos \theta_\odot)' = 1 + (\cos \theta_\odot - 1)(1 + \Delta_\theta). \quad (9.9)$$

The constraint on this parameter is 0.0 ± 0.11 . This makes for a total of twelve floating systematic parameters: $\Delta_S^{(0)}$, $\Delta_S^{(1)}$, Δ_R , and Δ_θ for each of the three phases.

Other parameters that affect PDF shapes are scanned, meaning that fits are run with the parameter fixed to different values, and the resulting maximum likelihoods used to determine the best-fit value and uncertainties. It is assumed that these parameters are uncorrelated among themselves (they may still be correlated between phases), and errors will be overestimated to the extent that nontrivial correlations exist. This is the same set of assumptions as was made for the counting analysis; the internally-correlated systematics there (low-energy signal extraction scalings and energy scale and resolution) are floated in the fit. For the signal extraction fit, the set is restricted to systematic parameters which were found in the counting analysis to have an impact on the signal or background of $\gtrsim 1\%$ in at least one phase. Specifically, the impact of the instrumental background cut sacrifice, live time, vertex resolution, ${}^8\text{B}$ ν spectrum shape, θ_{13} , and atmospheric neutrino-related systematics on the extracted uncertainties are assumed to be negligible. This leaves θ_{12} , Δm_{21}^2 , and the CC $\nu - d$ cross section as scanned parameters.

9.7 Ensemble Testing

As outlined in Section 9.4.3, the enforcement of a boundary at zero for rate parameters results in distributions for both the parameters and the likelihood function which are highly non-Gaussian. It is important to quantify the effect this has on the extracted confidence intervals, since clearly any claims of observing a nonzero *hep* flux must not be attributable to biases in the fit method. Neither bias nor pull distributions are good metrics for the *hep* flux fits near the boundary at zero, since the parameter distribution is one-sided and the “ 1σ ” errors not well-defined. Instead, I verify that the coverage is correct, using an ensemble of fits to fake data with the level of *hep* signal varying from a small fraction to many times the SSM prediction. In this way, we can gain confidence that the intervals produced by the fit to the data have the correct coverage, irrespective of the true value of the flux.

The coverage is evaluated for both methods introduced in Section 9.4, the Bayesian (§9.4.1) and profile likelihood (§9.4.2) constructions. In each case, a large set of fits were performed, using fake data with the *hep* flux set to 0.01 – 10.0 times the SSM prediction (significantly larger than the 78% 5σ model uncertainty), and all normalizations Poisson-distributed. Since the fitter runs much more quickly without floating systematics (about 60 times faster, in the configuration for the full *hep* fit), a larger set of fits were performed with fixed systematics, and a smaller number with floating systematics to check consistency. The results of these tests are shown in Figure 9.6; both our Bayesian and profile likelihood methods for extracting intervals do indeed have correct coverage across the range of *hep* fluxes to which our measurement is sensitive.

While these tests demonstrate that both statistical approaches produce intervals with the correct frequentist coverage for the relevant range of true *hep* flux values, they do not imply a minimum signal level actually required for discovery; the determination of the sensitivity is discussed in the next section.

9.8 Results

With the procedure having been validated through ensemble tests with fake data, signal extraction fits were performed using the methods described in the previous sections on the same unblind 1/3 dataset that was used in the counting analysis.

9.8.1 Sensitivity

Along similar lines to the counting analysis, we may define a sensitivity for the signal extraction fit by performing fits to an ensemble of signal-free Monte Carlo datasets. For a set of fits performed for datasets with normalizations and systematic parameters sampled

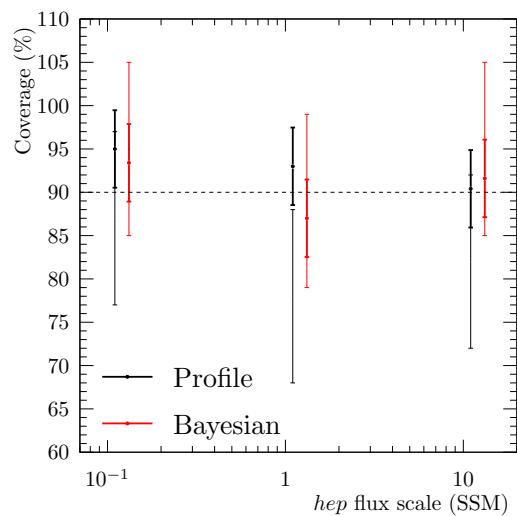


Figure 9.6: Coverage of 90% CL intervals in fake data ensemble tests using a variety of true *hep* fluxes. Error bars are statistical, due to the finite size of the ensemble. Thick lines are for fits with fixed systematics and thin lines for fits with floating systematics.

from their nominal distributions for the 1/3 dataset, the median upper limit is 2.27 times the BS05(OP) SSM, or

$$\langle \Phi_{hep}^{1/3} \rangle = 1.80 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL)}, \quad (9.10)$$

as calculated using the profile likelihood approach.

A sensitivity is computed in the same way for the full dataset, to estimate the improvement expected when the full statistics are analyzed. The median upper limit in this case is 1.16 times the BS05(OP) SSM prediction, or

$$\langle \Phi_{hep}^{\text{full}} \rangle = 9.21 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL)}. \quad (9.11)$$

This is within the 1σ errors for that model, suggesting that a first observation, if at low significance, may be achievable when the full dataset is unblinded.

9.8.2 1/3 Data Results

The maximum likelihood and uncertainties are calculated from an MCMC sample space containing 2.4×10^7 steps (not including burn-in steps), which provides $\mathcal{O}(5000)$ samples within the single-parameter 90% CL contours that define the profile likelihood intervals (see Section 9.4.2). One-dimensional projections of the best fit in the observable dimensions T_{eff} , β_{14} , and $\cos \theta_{\odot}$ are shown in Figure 9.8, and Table 9.2 summarizes the best-fit parameters and 90% CL uncertainties. With scanned systematics at their nominal values, the 90% confidence interval for the *hep* flux using the profile likelihood method includes zero, implying

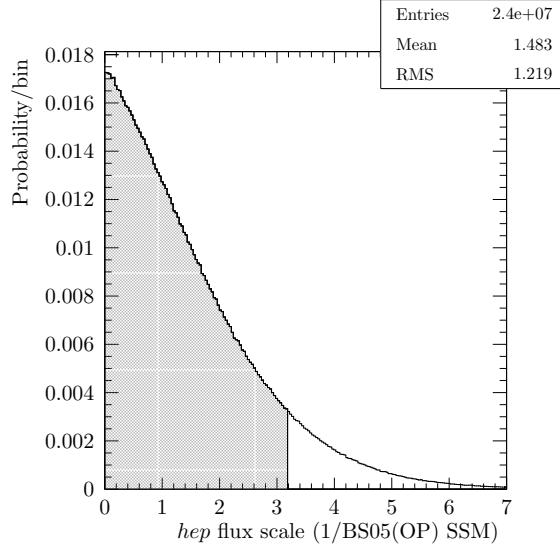


Figure 9.7: Marginalized posterior distribution for the *hep* flux, used to determine credible intervals. The 90% CL interval is highlighted in gray.

an upper limit of 2.38 times the BS(05)OP prediction, or

$$\Phi_{hep} < 1.90 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL)}, \quad (9.12)$$

which is consistent with the expected sensitivity. The corresponding limit in the Bayesian framework is 3.30 times the SSM, or

$$\Phi_{hep} < 2.61 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL)}, \quad (9.13)$$

with the marginalized posterior distribution for the *hep* flux shown in Figure 9.7.

Next, scanned systematics are evaluated by performing fits with these parameters shifted by $+1\sigma$ and -1σ relative to their means. Table 9.3 summarizes the results, including the

9.8 Results

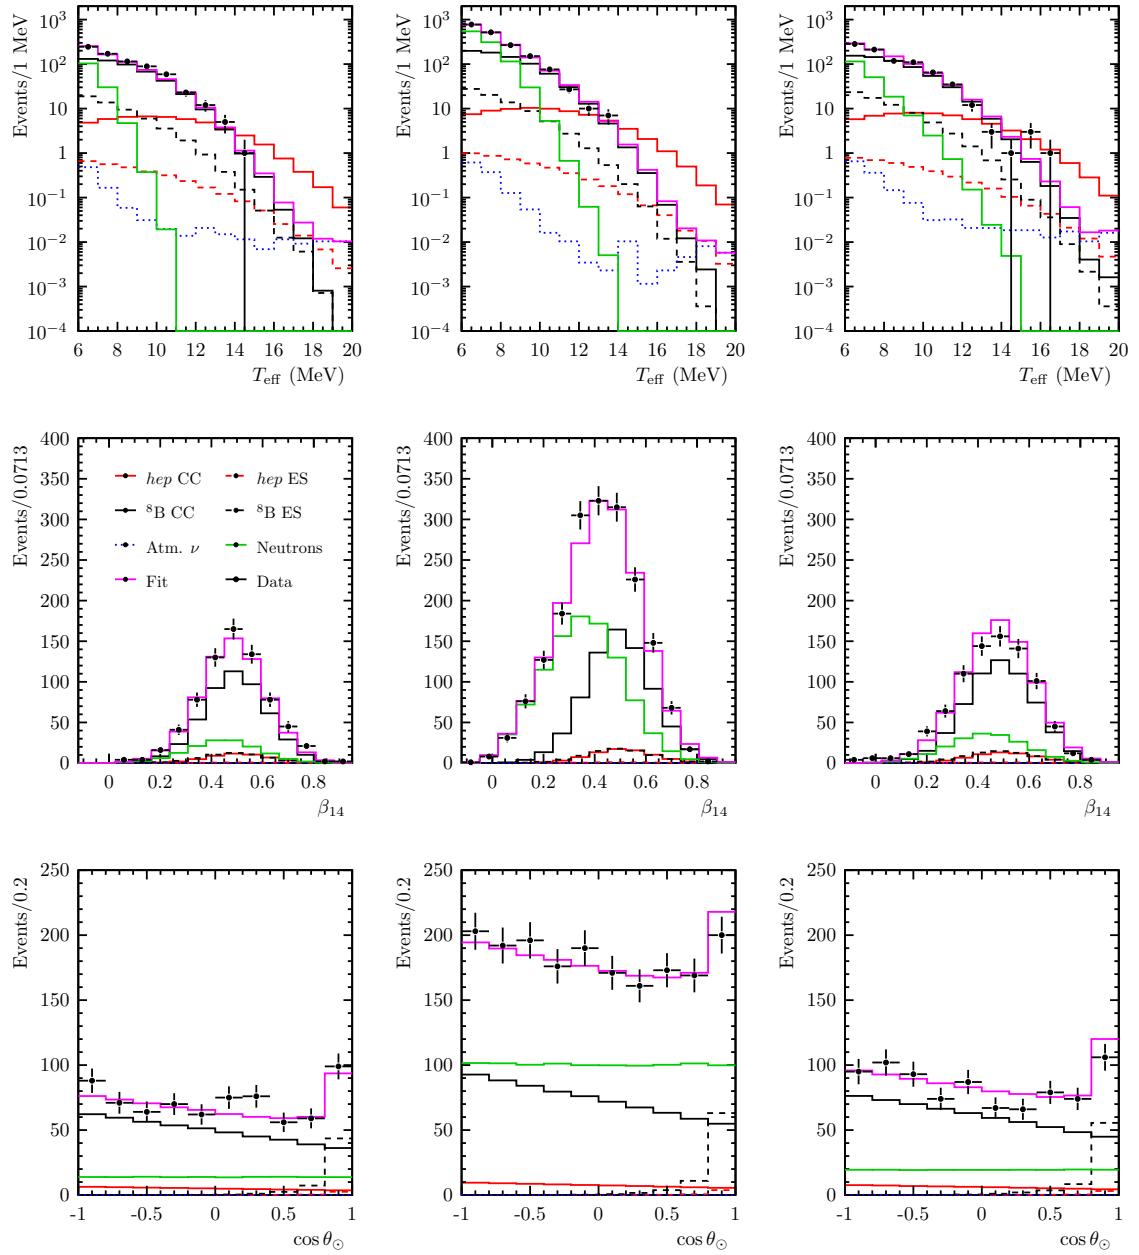


Figure 9.8: Projections of the best fit to the 1/3 dataset. The *hep* CC and ES signals are shown at the 90% CL upper limit. Phases I, II, and III are shown in the left, center, and right columns, respectively.

Parameter	Fit	Constraint	Correlated
${}^8\text{B}$ scale	$0.963^{+0.040}_{-0.041}$	—	✓
hep scale	$0.082 (< 2.034, 90\% \text{ CL})$	—	✓
Atm. ν scale	$1.015^{+0.189}_{-0.221}$	$\pm 15\%$	✓
Neutron scale I	$1.384^{+0.398}_{-0.204}$	—	
Neutron scale II	$1.526^{+0.114}_{-0.091}$	—	
Neutron scale III	$0.960^{+0.191}_{-0.188}$	—	
$\Delta_S^{(0)} \text{ I } /10^{-3}$	$0.855^{+10.302}_{-7.822}$	-5.20 ± 7.21	
$\Delta_S^{(1)} \text{ I } /10^{-3}$	$0.324^{+0.681}_{-0.419}$	0.44 ± 0.42	
$\Delta_S^{(0)} \text{ II } /10^{-3}$	$-4.954^{+7.554}_{-4.648}$	-0.01 ± 6.14	
$\Delta_S^{(1)} \text{ II } /10^{-3}$	$-0.333^{+0.455}_{-0.353}$	-0.16 ± 0.37	
$\Delta_S^{(0)} \text{ III } /10^{-3}$	$12.330^{+8.769}_{-4.866}$	-1.25 ± 10.20	
$\Delta_S^{(1)} \text{ III } /10^{-3}$	$0.282^{+0.336}_{-0.596}$	-0.16 ± 0.43	
$\Delta_R \text{ I } /10^{-2}$	$2.951^{+0.863}_{-2.809}$	1.83 ± 1.60	
$\Delta_R \text{ II } /10^{-2}$	$1.529^{+1.523}_{-1.268}$	2.38 ± 1.71	
$\Delta_R \text{ III } /10^{-2}$	$2.639^{+0.700}_{-2.181}$	1.61 ± 1.37	
$\Delta_\theta \text{ I}$	$-0.016^{+0.189}_{-0.123}$	0.0 ± 0.11	
$\Delta_\theta \text{ II}$	$-0.023^{+0.164}_{-0.098}$	0.0 ± 0.11	
$\Delta_\theta \text{ III}$	$0.037^{+0.126}_{-0.138}$	0.0 ± 0.11	

Table 9.2: Best-fit values with 90% CL profile likelihood errors and 1σ constraints for the parameters floated in the joint three-phase signal extraction fit to the 1/3 dataset, with scanned systematics at their nominal values. Flux scales here are relative to the BP2000 model prediction.

Parameter	Constraint	Best Fit	Shift	<i>hep</i> 90% CL UL	
				Profile	Bayesian
Nominal	—	—	—	2.39	3.30
$\tan^2 \theta_{12}$	0.437 ± 0.29	0.407	-1σ $+1\sigma$	3.81 2.24	4.04 3.59
$\Delta m_{21}^2 / 10^{-5}$	7.53 ± 0.18	7.51	-1σ $+1\sigma$	3.29 2.20	4.14 3.51
$\sigma(\text{CC } \nu - d)$	$\pm 1.2\%$	$+0.27\%$	-1σ $+1\sigma$	2.13 2.39	3.33 3.24

Table 9.3: Upper limits in units of the BS05(OP) SSM prediction, obtained with scanned systematics parameters individually varied in the joint three-phase signal extraction fit to the 1/3 dataset.

hep flux limits with each parameter shifted individually. Figure 9.9 shows the maximum likelihoods for the shifted fits, suggesting that the data is consistent with the prior constraints.

To derive a conservative combined limit for the profile likelihood case, the 90% CL MCMC sample subspaces for the nominal fit and for all fits with scanned systematics varied $\pm 1\sigma$ are combined, and the uncertainties are recalculated. This approach yields an upper limit of 3.8 times the BS(05)OP SSM prediction, or

$$\Phi_{hep} < 3.01 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL).} \quad (9.14)$$

For the Bayesian approach, the sample spaces are combined to form an average posterior distribution with which the credible intervals are calculated. The resulting limit is 3.9 times the SSM, or

$$\Phi_{hep} < 3.08 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1} \text{ (90\% CL).} \quad (9.15)$$

A more comprehensive approach would account for the differences in likelihood when the

systematic parameters are shifted. For example, we may sample systematics from their prior distributions, run fits with many sampled values, and combine the likelihood spaces to compute intervals accounting for the change in likelihood relative to the global minimum. This computationally-intensive approach will be pursued for the analysis of the full data set.

These limits represent a significant improvement over the results obtained in the counting analysis (Section 8.5). Moving beyond that one-bin analysis, we have used the PDF shape information across multiple bins, correlations in the fluxes across the phases of SNO, an improved, simultaneous measurement of the ${}^8\text{B}$ background rate, and the concurrent optimization of dominant systematic uncertainties along with the signal normalizations.

This limit is consistent with the previous work toward a fit for the *hep* flux [16]. Unlike that work, which indicated a nonzero flux at the $\sim 68\%$ confidence level, here no such excess is observed in the data. This difference may be due entirely to statistical fluctuations, as these analyses were performed using different 1/3 samples of the data. Issues with the past analysis may also contribute, such as a too-low constraint on the atmospheric neutrino background, the use of energy systematics extrapolated from low-energy calibrations, or perhaps a bias in the method (ensemble tests were performed only with large numbers of *hep* events, not in the small-signal regime of interest). While the lack of observation is perhaps a disappointing finding, the reanalysis of the energy response and atmospheric neutrino model, along with the validation of the fit method through ensemble testing, have been performed to ensure the limit is robust.

9.8.3 Goodness of Fit

The quality of the fit is assessed through a variety of metrics. First, there is the question of MCMC convergence: have we adequately sampled around the true global minimum? The time series for each parameter is shown in Figure 9.11, and indicate that the good mixing

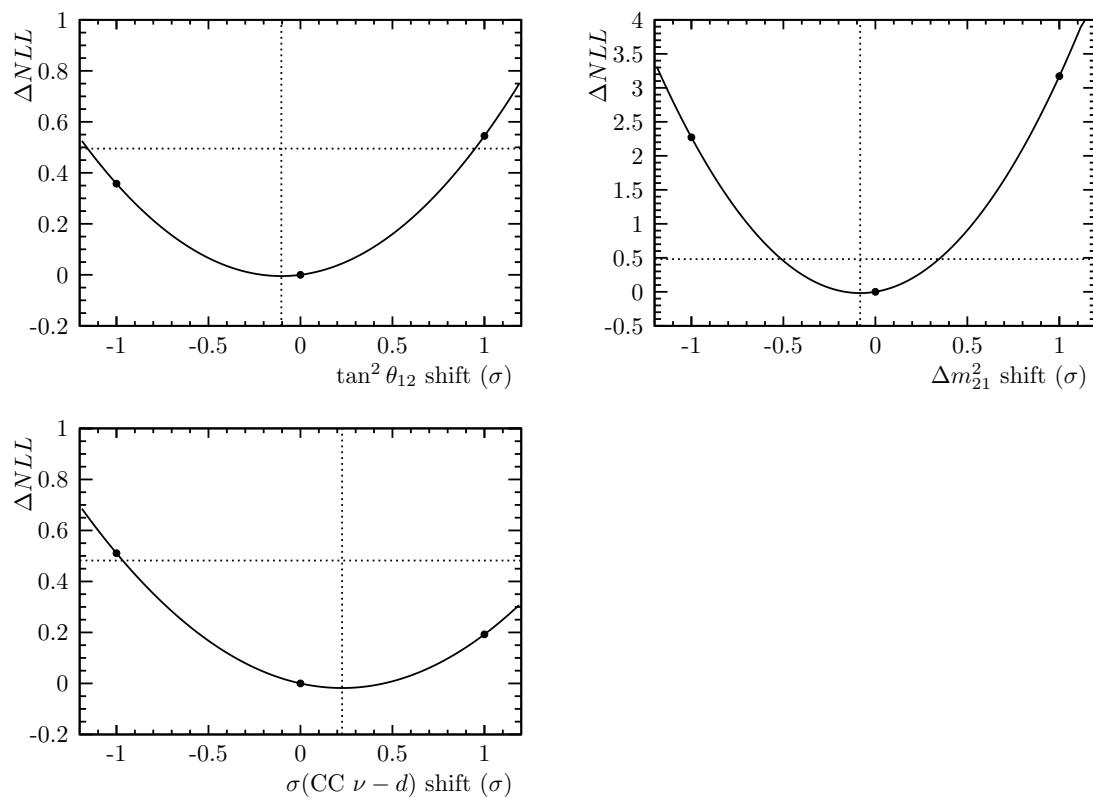


Figure 9.9: Fits to shifted systematic parameters.

9.8 Results

has been achieved, and this is verified by KS tests of time slices as described in Section 9.3.1.

Although a stationary distribution has been reached, the likelihood space alone cannot indicate whether the quality of the fit is good. For this the Pearson's χ^2 test is used, where data are binned for comparison to the (three-dimensional) best-fit histograms for each phase (see, e.g., Reference [108]). The expectation values in each bin are very small over much of the PDF domain, and so we are far from the asymptotic limit where the χ^2 statistic is expected to be distributed according to the χ^2 PDF, $\chi^2(k - 1 - n_{dof})$. Therefore, in order to calculate a p -value for the fit to data, I first construct this PDF, denoted $\xi(k)$, explicitly by sampling many independent datasets from the best-fit histograms and computing the standard χ^2 statistic; this distribution is shown in Figure 9.10. The p -value for a given fit, characterized by a χ^2 statistic a , is then given by

$$p = \int_a^\infty \xi(k) dk. \quad (9.16)$$

The χ^2 statistic for the 1/3 data fit is 5474, corresponding to a p -value of 25.8%, indicating that the data are fully consistent with the best-fit hypothesis.

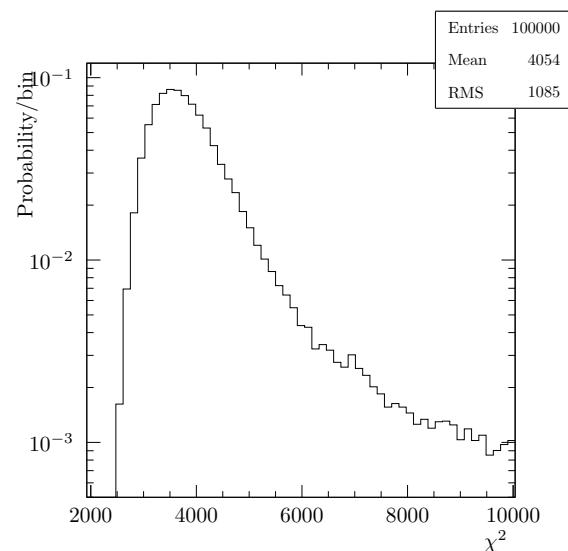


Figure 9.10: Monte Carlo-derived probability distribution for the Pearson χ^2 statistic.

9.8 Results

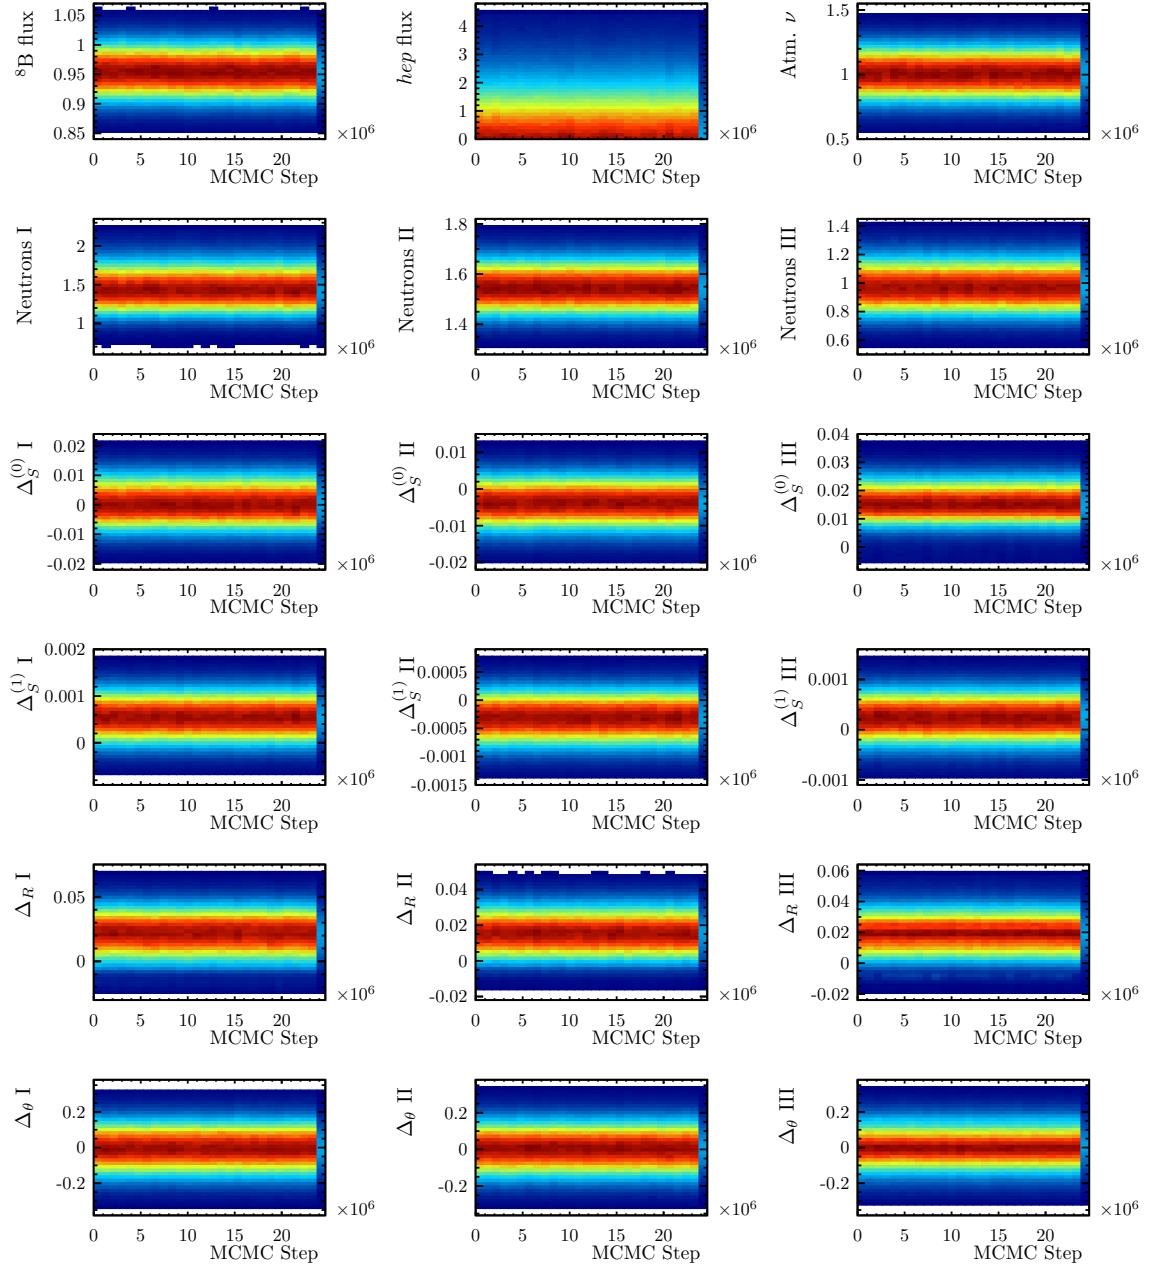


Figure 9.11: Parameter values as a function of time in the MCMC likelihood evaluation.

Chapter 10

Conclusions

Despite significant progress in terrestrial neutrino detectors over the past decades, the fluxes due to *hep* solar neutrinos and the diffuse supernova neutrino background remain unobserved. The *hep* flux is the last holdout from the *pp* chain, and observation will finally complete the validation of — or uncover issues with — the standard solar model. The DSNB spectrum holds valuable information about core-collapse supernova dynamics: the average number and temperature of emitted neutrinos, the particles which carry away most of the energy from these systems.

Data from all three operational phases of the Sudbury Neutrino Observatory experiment has been analyzed to search for *hep* solar neutrinos and the diffuse supernova neutrino background, using a dataset representing 1/3 of the total live time and equivalent to 0.82 kton years of exposure. I have performed this search with two different approaches: a counting analysis, wherein a sensitivity-optimized search region is defined in a space of cuts (*hep* and DSNB) and a multidimensional signal extraction fit which benefits from improved statistics and the use of PDF shapes (*hep* only). No evidence of the *hep* or DSNB flux was

observed at the 90% CL, however this is the most sensitive search performed to date, and a *hep* flux at or near the BS(05)OP standard solar model prediction is likely measurable in the full dataset. In support of this work, I have performed an independent re-calibration of the systematic uncertainties associated with energy reconstruction in the SNO detector, using data from a *pT* source, ⁸Li source, and Michel electrons, combined with previously published results using ¹⁶N at low energies.

For the counting analysis I have used a modified Feldman-Cousins approach where systematics are sampled with a Monte Carlo method. The sensitivity obtained for the *hep* flux at the 90% confidence level is 4.19 times the BS05(OP) SSM prediction, or $33.2 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1}$. With six events observed, an upper limit of 6.51 times the SSM is obtained, or $51.7 \times 10^3 \text{ cm}^{-2} \text{ s}^{-1}$. The corresponding sensitivity for the previously published Phase I-only analysis [15] is 4.34 times the SSM, and upper limit is set at 2.9 times the BP2000 SSM (2.47 times BS05(OP)) based on two events observed.

No events are observed in the DSNB search region; the upper limit is 61.6 times the $T = 6$ MeV Beacom and Strigari model [63], corresponding to a limit of $\Phi_{\nu_e}^{\text{DSNB}} < 41 \text{ cm}^{-2} \text{ s}^{-1}$ in the energy range $22.9 < T_{\text{eff}} < 36.9$ MeV.

A signal extraction fit for the *hep* flux was performed using an unbinned maximum likelihood approach implemented with a Markov Chain Monte Carlo. I have performed a simultaneous three-dimensional fit the full three-phase dataset, floating twelve systematic parameters, to extract the neutrino fluxes. Additional, much smaller systematic uncertainties were included via shift-and-refit approach. Parameter confidence intervals were calculated according to both a profile likelihood (frequentist) and a Bayesian approach. The signal extraction yields a 90% CL upper limit on the *hep* flux of $\Phi_{\text{hep}} < 1.90 \times 10^4 \text{ cm}^{-2} \text{ s}^{-1}$, (2.38 times the BS(05)OP SSM prediction) using a profile likelihood approach. This improves on previous SNO results, and suggests that a measurement (if at low significance) may be within reach with the full statistics.

In future work, these techniques will be applied to the full SNO dataset, in order to either make the first measurement of the *hep* flux or set a very strong limit that will likely stand for years to come. An improved DSNB limit will also provide useful input to supernova models. SNO has a unique capability to make these measurements, and it is crucial that we make these data available to the physics community.

Appendices

Appendix A

The SNO+ Experiment

The SNO+ experiment, currently under construction, will explore a broad physics program including nucleon decay, solar neutrinos, reactor antineutrinos and geoneutrinos, and neutrinoless double-beta decay ($0\nu\beta\beta$), with the $0\nu\beta\beta$ as the primary objective. This project, initiated in 2004 by a subset of the SNO collaboration led by M. Chen, reuses the SNO detector infrastructure, but replaces the D_2O target with Tellurium-loaded liquid scintillator. The modifications made to the SNO detector to transition to SNO+ are described in Section 5.7.

The unique capabilities of the SNO+ detector make it well-poised to make meaningful contributions in several areas of neutrino physics and astrophysics. The primary virtues of the detector are the large target volume, excellent shielding due to depth, and the ease with which the target material can be changed. The SNO+ collaboration will leverage the latter point by taking a phased approach to the experiment, exploring different physics along the way to a neutrino-less double beta decay measurement.

In the first phase, a water-filled detector will be used for an invisible nucleon decay

A.1 Neutrinoless Double Beta Decay

search, and also provide calibration data for comparison to SNO. Next, liquid scintillator will replace water in the active volume, yielding a configuration similar to Borexino [7]. This phase provides opportunities to study reactor $\bar{\nu}$ and potentially low-energy solar neutrinos, if backgrounds are sufficiently low. Scintillator-phase background measurements provide constraints to a third phase, when Te is added to the scintillator for the neutrinoless double beta decay search.

SNO+ will be the first experiment to load Te into liquid scintillator, making use of a novel technique [123]. Key advantages to this approach are that very large target mass is possible with further R&D to improve optics, and that switching isotope is possible in order to confirm an observed signal. Liquid scintillator detectors have received criticism due to their relatively poor energy resolution in a field historically dominated by small, high-precision detectors (Ge, bolometers, etc.). The collaboration aims to demonstrate both the soundness and the scalability of the SNO+ model: essentially, sensitivity scales as S/\sqrt{B} , hence more signal (large detector) is ultimately better than lower backgrounds (good resolution).

As of mid-2016, the SNO+ detector is currently nearing the end of the construction phase and due to begin commissioning soon; finalizing scintillator purification systems is the largest remaining task. Water filling is underway, suggesting a start to the first phase in late 2016. The next (scintillator) phase is scheduled to begin early 2017, and the third (double-beta decay) phase in the latter part of 2017. The nominal plan for this third phase calls for a five-year run time.

A.1 Neutrinoless Double Beta Decay

Neutrino oscillation experiments in recent decades have firmly established that the neutrino is massive. However, with no charges to distinguish ν from $\bar{\nu}$, it remains unknown whether the neutrino is described by a Dirac or Majorana field. The particle-antiparticle distinction

A.1 Neutrinoless Double Beta Decay

has historically been inferred from the other participants in a interactions, with an assumption of total lepton number (L) conservation. However, the motivation for this assumption is questionable: L conservation is an accidental symmetry of the Standard Model, and lepton flavor conservation is already known to be violated by neutrino oscillations. We may turn the question around: if you run faster than a (massive) ν_L , what about it *isn't* an antineutrino?

The current experimental evidence allows a Majorana neutrino, for which $\nu = \nu^c$ and $\nu_R = (\nu_L)^c$; this neutrino is its own antiparticle. This symmetry permits a variation of double-beta decay ($2\nu\beta\beta$) wherein the (virtual) Majorana neutrinos effectively undergo $\nu - \bar{\nu}$ oscillations and are absorbed internally, resulting in a three-body rather than a five-body decay; this is known as neutrinoless double beta decay ($0\nu\beta\beta$). This decay has a clear experimental signature in a scintillator detector where the individual electron tracks are not resolved: the observed sum of the energies of the outgoing electrons is a peak rather than a continuum, as shown in Figure A.1. $0\nu\beta\beta$ is not the only probe of the Majorana nature of the neutrino, but it benefits from a level of model-independence: the Schechter-Valle theorem [124] guarantees that any such $\Delta L = 2$ decay can be re-expressed as neutrino-antineutrino oscillation and thus implies a Majorana mass term in the Standard Model. A simple realization of this, most commonly used when discussing $0\nu\beta\beta$, is light neutrino exchange, illustrated in Figure A.2.

The parameter being probed in $0\nu\beta\beta$ experiments is the effective Majorana neutrino mass, to which the lifetime is inversely proportional. This is a coherent sum of neutrino eigenstate masses (including two new Majorana CP phases ϕ_1 and ϕ_2 in the mixing matrix):

$$(t_{1/2}^{0\nu})^{-1} = g_A^4 G_{0\nu}^{(0)} |M_{0\nu}|^2 \left| \frac{\langle m_{\beta\beta} \rangle}{m_e} \right|^2, \langle m_{\beta\beta} \rangle = \sum_i U_{ei}^2 m_i \quad (\text{A.1})$$

where G and M are the theoretical phase space factor and matrix element for the decay, and

A.1 Neutrinoless Double Beta Decay

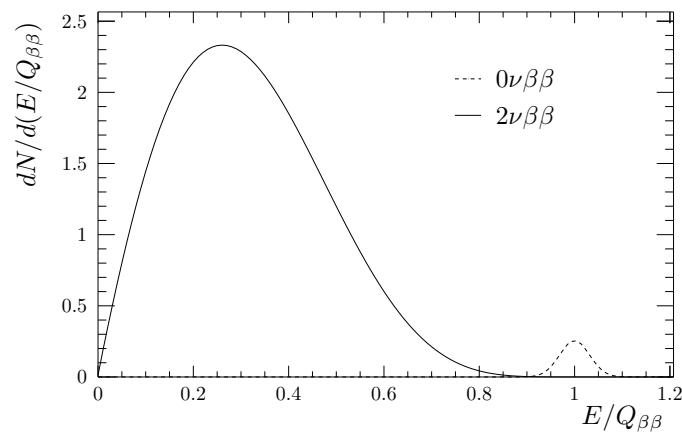


Figure A.1: Double-beta decay spectra, including a few-percent energy resolution.

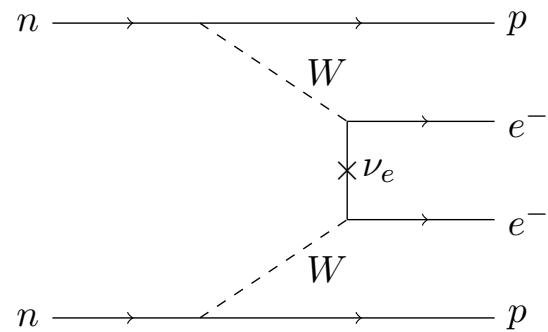


Figure A.2: Light Majorana neutrino exchange, one possible mechanism for $0\nu\beta\beta$.

A.1 Neutrinoless Double Beta Decay

the normalization of the axial vector coupling, g_A , has been factored out of G to show the strong dependence of the rate. This is of great interest since loop-level corrections involving axial vector currents can lead to an effective tree-level g_A that is smaller than the 1.269 appropriate for free nucleons, leading to a systematic underestimate of $m_{\beta\beta}$ limits when translating from $t_{1/2}$ limits. Figure A.3 shows the parameter landscape for $0\nu\beta\beta$ detection in a Vissani-Strumia plot, with $m_{\beta\beta}$ as a function of the lightest neutrino mass (ν_l), which depends on the neutrino mass ordering. In the normal ordering case ($m_1 < m_2 < m_3$), there are some values of ν_l and the Majorana phases which lead to a precise cancellation; in this experimentally unfortunate case, $m_{\beta\beta}$ would have an unobservably small value, despite Majorana neutrinos being realized in nature. In the case of the inverted mass ordering, $0\nu\beta\beta$ should be observable by next-generation experiments currently being planned. It may also be the case that $m_{\beta\beta}$ lies in the degenerate region, in which case it may be accessible to current-generation experiments. Finally, it is important to keep in mind that this parameter space is model-dependent; the landscape is very different, for example, if there exist one or more sterile neutrino states.

Figure A.3 also shows limits from cosmological structure observations by Planck [64], which constrain the sum of neutrino masses, the best limits from tritium beta decay endpoint measurements (the combined limit of the Mainz and Troitsk experiments [28]) which measure

$$m^2(\nu_e) = \sum_i |U_{ei}|^2 m_i^2, \quad (\text{A.2})$$

and a claim of observation of $0\nu\beta\beta$ in ^{76}Ge [125], with nominal ($g_A = 1.269$) and maximal quenching.

The Standard Model $2\nu\beta\beta$ process has been observed in a variety of nuclei for which competing processes are forbidden, with lifetimes on the order of 10^{21} years and energies around 3 MeV. The natural abundance of the $\beta\beta$ isotopes range from < 1% to about 35%, with a variety of matrix elements M , phase space factors G , and Q values. This leaves an

A.1 Neutrinoless Double Beta Decay

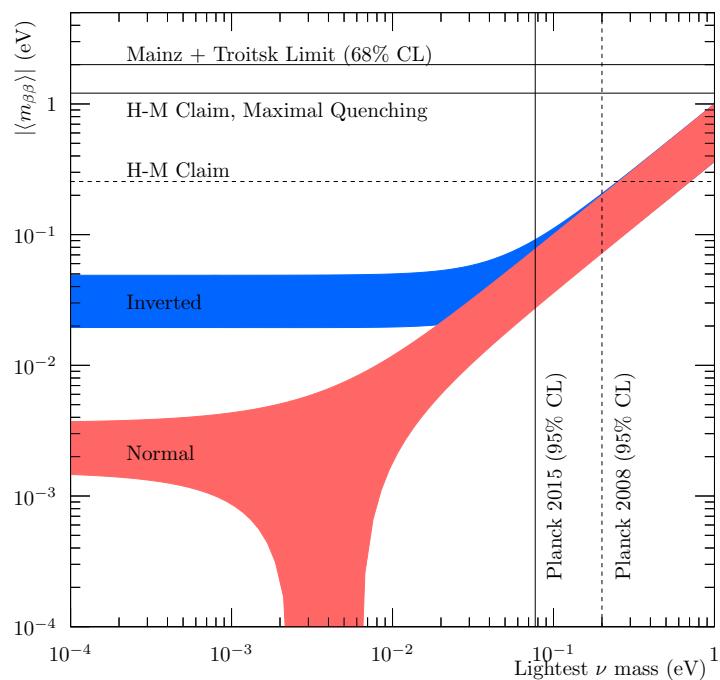


Figure A.3: A Vissani-Strumia plot showing the parameter space for $0\nu\beta\beta$ in the canonical three-neutrino mixing model.

array of possible isotopes with which to study $0\nu\beta\beta$, each with particular advantages and drawbacks, with no clear best choice [126]. ^{130}Te , the choice for the SNO+ experiment, is favorable due to a high natural abundance (34.5%), reducing need for costly enrichment, as well as a relatively large matrix element. However, the 2.528 MeV endpoint coincides with several backgrounds including cosmogenically-activated ^{60}Co and a 2.6 MeV γ s from U- and Th-chain decays in detector materials, which produce electrons through Compton scattering.

A.1.1 Current Results

A number of experiments are currently performing searches or due to begin taking data soon, and have adopted a wide variety of approaches and candidate isotopes.

Since 2006 there has been a highly controversial claim of detection in ^{76}Ge by a subset of the Heidelberg-Moscow experiment, with a measured half-life of $T_{1/2} = (2.23^{+0.44}_{-0.31}) \times 10^{25}$ y [125]. The authors have applied a variety of analysis improvements and claim a very high significance ($> 6\sigma$), but the experiment has received significant criticism from the rest of the community; see, e.g., Reference [127]. In the mean time, significant experimental effort has gone toward confirming or excluding the result, particularly using the same isotope. In 2013, the ^{76}Ge -based GERDA experiment published a null result with an upper limit of $T_{1/2} > 2.1 \times 10^{25}$ y (90% CL), which is combined with exclusions from other Ge experiments (Heidelberg-Moscow and IGEX) to reach $T_{1/2} > 3.0 \times 10^{25}$ y (90% CL), excluding the claim with high significance [128]. The authors of the claim have offered a rebuttal of this conclusion [129]. Very recently (July 8, 2016) the GERDA collaboration presented preliminary results combining new data from an upgraded Phase II run with the existing data; the ^{76}Ge half-life limit is $T_{1/2} > 5.2 \times 10^{25}$ y, with a sensitivity of 4.0×10^{25} y, both at the 90% CL.¹⁸

¹⁸These preliminary results are not yet published, and were presented at the XXVII International Conference on Neutrino Physics and Astrophysics (Neutrino 2016).

A.1 Neutrinoless Double Beta Decay

Meanwhile, the rest of the field of $0\nu\beta\beta$ experiments continue to improve sensitivity. The NEMO experiment uses thin foils with a number of $0\nu\beta\beta$ candidate isotopes in a tracking detector; although the target mass is low, NEMO is able to resolve the individual β kinematics. The NEMO-3 detector has set limits for a wide array of isotopes, including $T_{1/2} > 1.3 \times 10^{23}$ y (90% CL) for ^{130}Te [130].

The Enriched Xenon Observatory program has run EXO-200, a liquid Xenon time projection chamber, to search for $0\nu\beta\beta$ in ^{136}Xe , setting a limit of $T_{1/2} > 1.1 \times 10^{25}$ y (90% CL) [131]. The collaboration has also measured the $2\nu\beta\beta$ mode [132, 133].

The KamLAND-Zen experiment also uses ^{136}Xe , but dissolved in liquid scintillator, inside a thin containment balloon at the center of the KamLAND detector. The first phase produced a limit of $T_{1/2} > 1.9 \times 10^{25}$ y, or $T_{1/2} > 3.4 \times 10^{25}$ y (90% CL) when data are combined with EXO-200 [134]. This measurement was unfortunately limited by an unexpected contaminant near the Q value, believed to be ^{110m}Ag due to fallout from the Fukushima nuclear incident in March 2011, introduced during construction of the inner balloon. After a substantial background reduction campaign, the collaboration has recently announced¹⁹ a dramatically improved limit of $T_{1/2} > 1.1 \times 10^{26}$ y (90% CL) [135].

The Cryogenic Underground Observatory for Rare Events (CUORE) program is using ^{130}Te , fabricated into TeO_2 crystals operated cryogenically as bolometers. The Cuoricino prototype operated at LNGS from 2003–2008 set a limit of $T_{1/2} > 2.8 \times 10^{24}$ y (90% CL) [136], while in 2015 the CUORE-0 detector (a representative fraction of what will become the full CUORE detector) improved this to $T_{1/2} > 4.0 \times 10^{24}$ y (90% CL), the strongest limit for ^{130}Te thus far [137].

The experiments using isotopes other than ^{76}Ge may still test the claim of observation in ^{76}Ge , but the constraints are weakened by the uncertainties in the matrix element used to convert from a half-life (which is different for every isotope) to the universal $m_{\beta\beta}$. Even so, the published joint KamLAND-Zen/EXO-200 results exclude the claim regardless of the

¹⁹At the time of writing, this paper has not yet been published.

matrix element calculation. Furthermore, cosmological constraints on the lightest neutrino mass due to the Planck experiment [64] independently exclude the region of the degenerate parameter space required for the claim at high significance.

In light of this evidence, the ${}^{76}\text{Ge}$ claim is effectively ruled out, and so there is no longer a particular sensitivity target in mind. Therefore, all of these experiments — and more, including SNO+ and MAJORANA [138] — are aggressively pursuing new technologies that have a clear path forward to covering the entire inverted hierarchy region and beyond.

A.2 Backgrounds

With the SNO+ detector still under construction, analysis relies heavily on detector Monte Carlo. To this end, we have developed a sophisticated detector simulation based on GEANT4 [139, 140] which includes geometry details and a microphysical optical model with photon tracking.²⁰ Most decays, including double-beta decays, are generated with the Decay0 generator created by V. Tretyak and ported to C++ by A. Bialek of SNO+. We model all expected backgrounds at all locations in the detector, as described below. Microphysics, in particular the scintillation and PMT model, is verified with benchtop experiments, and the detector modeling has been checked against both the SNO Monte Carlo and SNO data.

Internal backgrounds originate in the scintillator, the Te, and the agents used to load the Te into the scintillator. These include U- and Th-chain contamination, cosmogenically-activated isotopes, and $2\nu\beta\beta$ decays. We also group elastic scatters of ${}^8\text{B}$ solar neutrinos into this category, as well as n -capture γ s. The estimated rates of these backgrounds come from a variety of sources: scintillator is assumed to purified to Borexino levels [142], ${}^8\text{B}$ ES from the global best fit to solar and reactor neutrino measurements, and $2\nu\beta\beta$ from NEMO-3 [130]. Based on spike tests, it is believed that cosmogenically-activated isotopes

²⁰The author has led the release of an open-source variant of this software, known as RAT-PAC (RAT is an Analysis Tool, Plus Additional Codes) [141].

resulting from the exposure of the Te to cosmic rays during its time on the Earth’s surface will be reduced to negligible levels. A detailed analysis of the cosmogenic activation of Te for SNO+ has been performed by Lozza and Petzoldt [143].

External backgrounds originate outside the target volume; these are predominantly U- and Th- chain decays in the detector materials. Rates of these decays are estimated based on the target levels [144] and fit results [46] from SNO.

Some backgrounds have complex distributions in time and space: for example, diffusion of Rn from the cover gas into the scintillator, and leaching of Rn daughters from the acrylic vessel bulk into the scintillator. These processes are modeled in an average way, with rates estimated by *ex situ* assays.

A final background class, *instrumental* backgrounds are those due to noise in the detector electronics, or light generated in detector components rather than in decays and scatters. SNO+ benefits from the experience of SNO in identifying these backgrounds, but much work remains to be done in developing cuts appropriate for a scintillator-filled detector. These backgrounds are not modeled: it is assumed that they will be cut effectively by future analysis.

A full account of the SNO+ background model is given in Reference [56].

A.3 Counting Analysis

The measurement of $0\nu\beta\beta$ in SNO+ is in many ways similar to the *hep* and DSNB measurement in SNO described in the preceding chapters, despite the different physics involved. Both are rare-process searches at the endpoint of a steeply-falling background spectrum (^{130}Te $2\nu\beta\beta$ and ^8B solar neutrinos, respectively) and so the analysis approaches are quite similar. A counting analysis provides a straightforward and robust means of deriving a sensitivity comparable to other experiments, and a more sophisticated signal extraction fit to be performed in the future will extract the most information out of the available statistics.

A.3.1 Sensitivity

To estimate the sensitivity of the SNO+ experiment to $0\nu\beta\beta$, I employ the background fluctuation sensitivity introduced in Section 8.4; that is, I compute the number of background events that would be observed in a background fluctuation of a given significance, take that as a limit on the smallest observable signal, and translate to a limit on the decay lifetime and effective Majorana neutrino mass. An oft-quoted formula for the half-life sensitivity is [136]

$$\widehat{T_{1/2}^{0\nu}}(n_\sigma) = \frac{\ln 2}{n_\sigma} \frac{N_A \cdot a \cdot \eta \cdot \epsilon}{W} \sqrt{\frac{M \cdot t}{b \cdot \delta E}} \cdot f(\delta E) \quad (\text{A.3})$$

where n_σ is the desired confidence level in number of Gaussian σ , N_A Avogadro's number, a the isotopic abundance of $0\nu\beta\beta$ isotope, η the stoichiometric fraction of isotope, ϵ the detector efficiency, W the molecular weight of the active mass, M the active mass, t run time, b specific background (counts per unit energy-mass-time), δE the energy cut window, $f = \text{erf}(\delta E \sqrt{\ln 2} / \Delta E)$ the fraction of signal events inside the cut window, where ΔE is the FWHM energy resolution at the endpoint.

This expression, however, makes a series of assumptions that are not valid for SNO+. First, it assumed that the number of background counts is large enough that a Gaussian distribution models the Poisson uncertainty accurately, $N_b \gtrsim 25$. This is particularly dangerous because a Gaussian approximation will under-cover (set ‘too good’ a limit) for low N_b . Second, it is assumed that the signal rate and background rate scale together. This is the case for many detectors such as CUORE [137] and GERDA [128], where the detection medium itself is made from the candidate isotope, but in SNO+, the external backgrounds, the internal backgrounds intrinsic to the LAB-PPO scintillator, and the ^8B solar neutrino elastic scattering background do not scale with Te isotope mass.

For these reasons I adopt a more exact procedure based on the Feldman-Cousins unified approach [110]. Feldman-Cousin upper limits (FC) are derived for the desired confidence

A.3 Counting Analysis

level α assuming that a number of events N are observed that is Poisson-distributed around the Monte Carlo-derived background expectation b :

$$\widehat{T_{1/2}^{0\nu}}(\alpha) = \left\langle \frac{N_{\text{iso}} \cdot \epsilon \cdot t \cdot \ln 2}{\text{FC}(N, b; \alpha)} \right\rangle_{\{N=\text{Pois}(b)\}} \quad (\text{A.4})$$

where the average is over an ensemble of pseudo-experiments, N_{iso} is the number of ^{130}Te atoms inside the active (fiducial) volume, and ϵ is the signal detection efficiency (a generalization of $f(\delta E)$). This approach makes no assumptions about the relationship between target and detector mass, is valid in the Poisson limit, and limits may be compared to other Feldman-Cousins intervals.

While the above method is used for the baseline sensitivity calculation, note that all of these approaches assume not only that background targets are reached, but that the mean background rates are known exactly. This is certainly not the case for all backgrounds, and so we require a way to incorporate that uncertainty into our limits. Possibilities include, for example, a sampling approach similar to that used in Chapter 8, a hybrid Bayesian approach that modifies confidence belts by marginalizing $P(n; \mu)$ over the background distribution, or direct frequentist approaches. The sensitivities presented in the following sections, however, follow the basic approach of Equation A.4 with fixed b .

A.3.2 Reconstruction

Reconstruction algorithms with performance rivaling or exceeding those used in SNO are under development for SNO+. For the present analysis, however, I use Monte Carlo truth event positions and energies, and apply an analytical energy response model, convolving the true energy with a Gaussian kernel of width $\sigma_E = \sqrt{E \text{ [MeV]}/Y}$ where Y is the light yield in $N_{\text{hits}}/\text{MeV}$. An exception to this rule is external backgrounds: here both reconstructed

A.3 Counting Analysis

positions and energies are used. The SNO+ fitter is a 1D maximum likelihood fit, where the observable is the set of time-of-flight corrected PMT hit times (time residuals).

Due to quenching, the mean energy deposited by $0\nu\beta\beta$ events is less than $Q_{\beta\beta}$ (2.528 MeV) by about 50 keV, according to the simulation model. I define an effective kinetic energy unit $T_{\beta\beta}$ such that the mean visible energy for simulated $0\nu\beta\beta$ events sits at the Q value. It is difficult to construct such a unit rigorously, as it in principle involves an integral over the two-particle phase space, so for the purposes of this analysis we achieve a similar effect by simply multiplying all energies by a scale factor of 1.021.

A.3.3 Cuts

A number of event selection cuts are applied to reduce background and maximize the sensitivity to $m_{\beta\beta}$. It is assumed that these cuts have 100% signal efficiency (no sacrifice); this will be revisited once cuts are finalized and have been applied simultaneously to a single dataset.

BiPo Tagging Coincidence tagging is used to reduce the background due to $^{214}\text{Bi} - ^{214}\text{Po}$ and $^{212}\text{Bi} - ^{212}\text{Po}$ decays which occur in quick succession, the 212 family with a 299 ns half-life and 214 with 164 μs . We assume that we can tag and remove all events where the two decays fall into separate 400 ns trigger windows.²¹ In the case where decays “pile up” in the same window, techniques such as likelihood ratio tests using the time residuals provide discrimination up to the point where the decays are simultaneous within the resolution of the PMTs. Based on this limit, a factor of 50 rejection is assumed for these in-window coincidences. It is possible that particle ID afforded by the different α and β scintillation timing profiles in LAB-PPO scintillator will further aid in rejecting these backgrounds.

²¹The capabilities of the new analog trigger system developed by the author, the MTC/A+, may help to identify these coincidences; see Section 5.7.1.

External Backgrounds A likelihood ratio-based technique for identifying external ^{208}Tl events developed by Coulter [145] and independently verified by the author reduces external backgrounds by a factor of two. This method uses the distortion in the time profile of hit PMTs introduced when energy is promptly deposited in the AV or external water, before a γ ray Compton scatters in the scintillator.

(α, n) Backgrounds Coincidence tagging is also used to reduce the background due to (α, n) reactions, which can result in a 2.22 MeV capture γ which may contribute background in the signal region of interest. These tags cuts have been tuned to eliminate $> 99.6\%$ of prompt α events and 90% of delayed n events with minimal signal loss [56].

A.3.4 Optimizing the Signal Region of Interest

The $0\nu\beta\beta$ region of interest is defined by a restricted fiducial volume and energy window, which are set to maximize the sensitivity. It is common in the field — which is dominated by detectors with much better energy resolution — to use an energy window of 1σ or FWHM around the $0\nu\beta\beta$ signal peak to define the energy window. For SNO+, however, this is not the optimal choice, due to the presence of the steeply-falling $2\nu\beta\beta$ spectrum on the low-energy side. In terms of volume, we must balance signal acceptance with contamination due to external backgrounds, which have a roughly exponential radial profile. These issues are coupled: if we accept much more background by lowering the energy threshold, externals are subdominant and it makes sense to expand the fiducial volume slightly. Therefore, we have a two-dimensional optimization problem.

In order to address this, I performed a grid scan over a range of fiducial volume cut values and energy thresholds, and calculated an approximate sensitivity for each point in the parameter space, shown in Figure A.4. The best sensitivity is achieved for a radius cut around 3.6 m, which is rounded down to 3.5 to be more robust against uncertainties in the external background model, and an energy window of $-0.5\sigma \rightarrow 1.5\sigma$ around the

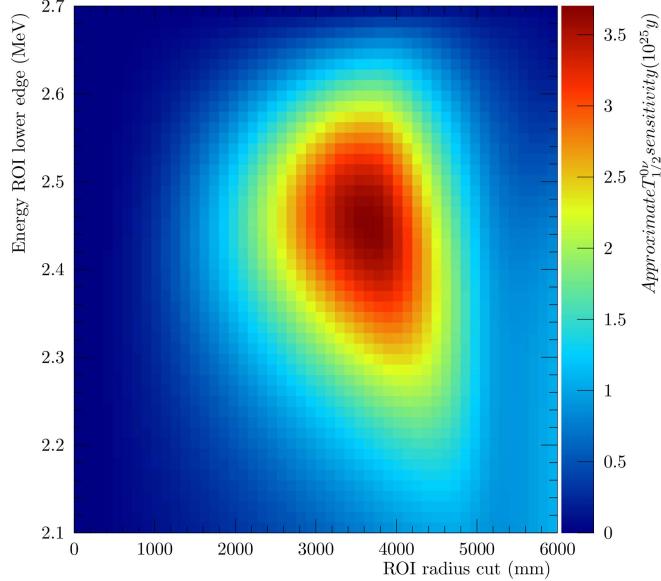


Figure A.4: Grid scan of the joint fiducial volume/energy range parameter space to optimize the signal ROI in the SNO+ counting experiment. Sensitivity is approximate only.

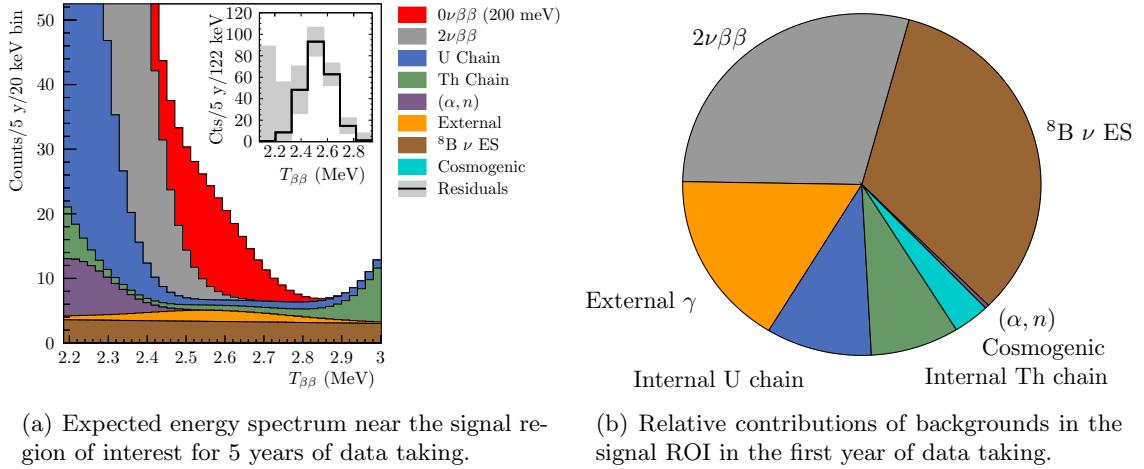
Gaussian $0\nu\beta\beta$ peak. The volume cut accepts 20% of signal events (the entire AV is filled with Te-loaded scintillator) and the efficiency of the energy cut is 0.625.

A.3.5 Results

Figure A.5 shows expected background distributions for the first year of running, assuming a scintillator cocktail with 0.3% ^{nat}Te by weight loaded using the surfactant PRS and a secondary fluor bis-MSB²² [123], with a fiducial volume cut at 3.5 m applied. The number

²²Since the time of this analysis, which provides the sensitivity estimates quoted in Reference [56], the SNO+ collaboration has pursued a new loading method based on complexing Te with 1,3-butanediol. The surfactant method suffers from (a) the higher absorption at low wavelengths, necessitating the introduction of a secondary fluor, and (b) the cosmogenic activation of sulfur in PRS precursor LAS, leading to very high

A.3 Counting Analysis



(a) Expected energy spectrum near the signal region of interest for 5 years of data taking.

(b) Relative contributions of backgrounds in the signal ROI in the first year of data taking.

Figure A.5: Expected backgrounds with $r < 3.5$ m, shown with an example signal for a Majorana neutrino with $m_{\beta\beta} = 200$ meV.

of counts expected in the ROI for the first year of data taking are listed in Table A.1.

Background rates in subsequent years change slightly for two reasons: short-lived cosmogenic backgrounds decay away, leading to reduction in this class, while leaching of Rn daughters from the inner surface of the AV leads to an increase in the rate of (α, n) backgrounds (as well as ^{210}Po , which affects our ability to tag certain $\beta - \alpha$ coincidences). These effects are accounted for in the sensitivity estimates for multi-year running, given in Table A.2.

The IBM-2 matrix element using the Miller-Spencer approach [146] and $g_A = 1.269$ [147] are chosen for definiteness in the figures, but this is entirely arbitrary; indeed, there is about a factor of two uncertainty among different theoretical approaches, and so it is more correct to quote a range. Table A.3 presents $m_{\beta\beta}$ ranges for a variety of matrix element calculations, where I have accounted for different choices of g_A in the original references

expected backgrounds in the ROI even after purification. Significant effort is underway to demonstrate the optical properties and long-term stability of the Te-diol complex.

A.3 Counting Analysis

Signal	Counts	Signal	Counts
$2\nu\beta\beta$	6.3297	^{22}Na	6.25×10^{-3}
$^8\text{B} \nu$ ES	7.2619	^{26}Al	1.65×10^{-9}
$(\alpha, n) \gamma$	0.0632	^{42}K	6.95×10^{-5}
$^{13}\text{C} (\alpha, n)$	0.0178	^{44}Sc	1.90×10^{-5}
$^{13}\text{O} (\alpha, n)$	0.0001	^{46}Sc	5.50×10^{-5}
^{234m}Pa	0.2726	^{56}Co	1.11×10^{-6}
$^{214}\text{BiPo}$	0.5413	^{58}Co	1.47×10^{-7}
^{210}Tl	1.2665	^{60}Co	4.75×10^{-3}
^{228}Ac	0.0004	^{68}Ga	1.62×10^{-3}
$^{212}\text{BiPo}$	1.7429	^{82}Rb	1.58×10^{-4}
^{208}Tl	0.0042	^{84}Rb	1.54×10^{-5}
		^{88}Y	2.07×10^{-1}
AV ^{214}Bi	0.2558	^{90}Y	3.11×10^{-9}
AV ^{208}Tl	1.0118	^{102}Rh	5.77×10^{-5}
AV Inner Dust ^{214}Bi	0.0028	^{102m}Rh	1.44×10^{-4}
AV Inner Dust ^{208}Tl	0.0000	^{106}Rh	1.44×10^{-4}
AV Outer Dust ^{214}Bi	0.0150	^{110}Ag	1.60×10^{-5}
AV Outer Dust ^{208}Tl	0.3952	^{100m}Ag	1.84×10^{-2}
Water ^{214}Bi	0.1508	^{124}Sb	4.72×10^{-1}
Water ^{208}Tl	0.4422	^{126}Sb	2.93×10^{-5}
HD Ropes ^{214}Bi	0.0237	^{126m}Sb	4.30×10^{-7}
HD Ropes ^{208}Tl	1.0429		
PMT $\beta - \gamma$	0.2788		
Total			21.8304

Table A.1: Expected background counts in the ROI in the first year of SNO+ data taking.

A.3 Counting Analysis

	1 Year	5 Years	10 Years
Counts	21.830	106.826	213.108
FC Limit	8.191	17.224	24.937
$\widehat{T_{1/2}^{0\nu}}$ [y]	3.95×10^{25}	9.38×10^{25}	1.30×10^{26}
$m_{\beta\beta}$ [meV]	105.1	68.1	58.0

Table A.2: Expected sensitivity of the SNO+ experiment at the 90% confidence level. See note in the text about the assumptions regarding the matrix elements used for the $m_{\beta\beta}$ estimate.

	IBM-2 [146]	QRPA-Tü [148]	ISM [149]	pnQRPA [150]	EDF [151]
$M^{0\nu}$	$4.03 - 4.61$	$3.89 - 4.81$	$2.06 - 2.57$	3.94	4.98
Phase I $m_{\beta\beta}$	$59.6 - 68.2$	$57.1 - 70.7$	$107 - 133$	69.7	55.2
Phase II $m_{\beta\beta}$	$20.5 - 23.5$	$19.7 - 24.3$	$36.8 - 45.9$	24.0	19.0

Table A.3: Effective Majorana mass limits for a nominal five-year SNO+ run, for a variety of theoretical approaches to matrix element calculation.

and normalized all to $g_A = 1.269$. The overall range for $m_{\beta\beta}$ for a nominal five-year SNO+ run is $55.2 - 133$ meV.

Constraints and Discovery Potential It is common to extend the background fluctuation sensitivity to a higher significance (say 3σ) and consider this as the signal level at which a discovery of $0\nu\beta\beta$ could be claimed. This interpretation can be misleading, however, as it does not address a crucial difference between discovery and limit-setting: in practice, a positive claim requires a demonstration that the observed signal could not have been caused by anything else. Imagine for the sake of argument that the ^{60}Co cosmogenic background is in all observables identical the $0\nu\beta\beta$ signal. In computing the sensitivity, we assume a model where this background is reduced to negligible levels by purification and exclude it.

And indeed, if the purification target is reached, the SNO+ limit could be interpreted as being distributed around that sensitivity estimate.

Realistically, the efficacy of the purification is unknowable at the relevant scale. If an excess were observed (due in reality to an unexpectedly high ^{60}Co rate), it could be explained by a short $0\nu\beta\beta$ half-life, and hence the limit is weakened. It is clearly not possible to make a discovery claim in this case, regardless of the confidence level of the background-fluctuation sensitivity. For this reason, it is crucial to preserve experimental handles that break the degeneracy between signal and background in the observables, even at the risk of diminished sensitivity. Systematic uncertainties are similarly problematic: a small non-Gaussian tail in the energy resolution could easily mimic a signal excess, and also must be tightly constrained.

Unless every component of the background model can be independently constrained, either through conclusive *ex situ* tests or better yet, using sidebands of the data itself, is it not possible to claim discovery. The strongest statement that we can make is to construct one-sided upper limits based on the known components of the background model as a function of the observed event rate, which will result in substantially weakened limits if the expectation is exceeded.

A.3.6 Systematic Uncertainties

There are many sources of systematic uncertainty relevant for the SNO+ $0\nu\beta\beta$ search: uncertainties in the true mean rates of the backgrounds, uncertainties related to position and energy reconstruction (including scale, shift, resolution, nonlinearity, tails, and time variation), uncertainties in the model ($Q_{\beta\beta}$, solar neutrino mixing parameters, the ^8B solar neutrino spectrum, cross sections, and physics models in Geant4), the target mass and effective fiducial volume, and the finiteness of the Monte Carlo samples used to evaluate

the limits. Additional uncertainties, such as the sacrifice of instrumental background cuts, must await detector data before they can be quantified.

The constraints on individual background rates are derived from a variety of analyses. The normalizations of the two dominant backgrounds, solar ${}^8\text{B}$ ν elastic scattering (ES) and $2\nu\beta\beta$, are constrained by previous experiments (SNO and NEMO-3 [130], respectively), but all other constraints must be determined from SNO+ data. This will be accomplished through coincidence tagging (e.g. n and β/α), sideband fits (e.g. energy and radius), and potentially time series analysis (to constrain certain cosmogenic backgrounds).

Systematic uncertainties other than normalizations also have impact on the sensitivity and discovery potential for SNO+. I have investigated a variety of systematics, both in terms of the minimum required constraint to keep backgrounds subdominant to irreducible ${}^8\text{B}$ ν scattering and how well the data itself constrains the parameters. Energy systematics (scale, shift, and resolution) were extracted from sideband energy fits in the 1.5–2.2 MeV range, and other parameters have been propagated using Monte Carlo and event reweighting. The systematic parameters studied are listed in Table A.4. Per the discussion in A.3.1, it is crucial to further develop *in situ* constraints for cosmogenics and ${}^{60}\text{Co}$ in particular, and to revise all constraints once detector background data becomes available, in order to set a realistic sensitivity estimate.

A.4 SNO+ Phase II

Following the successful completion of the above-described $0\nu\beta\beta$ search, nominally with a five-year run time, the SNO+ detector will be upgraded for a second phase (SNO+ Phase II). The baseline plan for Phase II includes an increased Te loading of 3% and an upgrade of the PMT array to new, high-quantum efficiency photodetectors such as the Hamamatsu R5912-100. The increased loading has a negative impact on optical attenuation in the scintillator, however the improved photon detection efficiency of the upgraded PMT array is

Parameter	Description	Nominal	Variation	Reference
σ_E [keV]	Energy resolution	112.4	$^{+2.4}_{-2.2}$	
a_E	Energy scale	0.021	$^{+3.4}_{-3.6} \times 10^{-3}$	
A_E [keV]	Energy offset	0	$^{+3.3}_{-3.7}$	
$Q_{\beta\beta}$ [keV]	$2\nu\beta\beta$ endpoint	2527.518	± 0.013	[152]
$\tan^2 \theta_{12}$	Solar ν mixing	0.443	$^{+0.030}_{-0.025}$	[48]
$\sin^2 \theta_{13} [\times 10^{-2}]$	Solar ν mixing	2.49	$^{+0.20}_{-0.32}$	[48]
Δm_{21}^2 [10 $^{-5}$ eV 2]	Solar ν mixing	7.46	$^{+0.20}_{-0.19}$	[48]
$P(E_\nu)$	^8B shape	See ref.	See ref.	[85]

Table A.4: Systematic parameters and constraints.

expected to more than compensate, for a total detected light yield of 450 N_{hit} /MeV. Based on preliminary studies, this upgrade is expected to improve the sensitivity by almost an order of magnitude, potentially allowing SNO+ to cover a large portion of the parameter space for the inverted neutrino mass hierarchy. With higher loading, fiducialization becomes increasingly inefficient, both in terms of optics and cost: a large amount of the Te isotope is “wasted” in the non-active region at high radius, and absorbs more light than unloaded scintillator. For this reason there is motivation to pursue an inner containment vessel, such as a thin balloon, which will have loaded scintillator inside and pure LAB-PPO outside. A significant R&D effort is already underway to characterize the higher-concentration scintillator cocktails, the response of a HQE PMT array, and the optimization of a possible containment balloon.

I have applied the counting analysis developed above to make projections for the Phase II sensitivity, assuming an increase to 3% loading and 450 N_{hit} /MeV. This involves a re-optimization of the volume/energy region of interest, which results in an expanded fiducial volume, out to 3.9 m in radius, while the ($-0.5\sigma \rightarrow 1.5\sigma$) energy ROI remains near optimal. Table A.5 shows how the assumptions, background rates, and sensitivities scale in moving from Phase I to Phase II.

	Phase I	Phase II
Fiducial volume [m]	3.5	3.9
Active isotope [kg]	160	2214
Light yield [N _{hits} /MeV]	200	450
Energy resolution σ at $Q_{\beta\beta}$ [keV]	112	75
Energy ROI	$(Q_{\beta\beta} - 0.5\sigma, Q_{\beta\beta} + 1.5\sigma)$	
$2\nu\beta\beta$	31.6	42.3
^8B ν ES	36.3	33.6
(α, n)	0.84	0.04
U chain	10.4	49.6
Th chain	8.74	49.6
Cosmogenic	0.79	6.37
External	18.1	63.7
Total counts	107	245
$T_{1/2}^{0\nu}$ (90% CL) [y]	9.4×10^{25}	8.4×10^{26}
$m_{\beta\beta}$ (90% CL) [meV]	68	23

Table A.5: Projections for SNO+ Phases I and II compared, each for a nominal five-year live time.

If the background and light yield targets for Phase II are achieved, a five-year measurement will have sensitivity at the 90% confidence level throughout the vast majority of the inverted mass hierarchy region of the $m_{\beta\beta}$ parameter space, with a limit of $T_{1/2}^{0\nu} > 8.4 \times 10^{26}$ y, or $m_{\beta\beta}$ sensitivity in the range 19.0 – 45.9 meV, where the range is due to the spread in matrix element calculations, as described in Section A.3.5.

A.5 Conclusions

SNO+ is a kilotonne-scale liquid scintillator detector that will follow on the success of the SNO experiment, reusing much the SNO detector infrastructure to perform a broad array

A.5 Conclusions

of measurements in neutrino physics. The primary goal for SNO+ is a sensitive search for $0\nu\beta\beta$ in ^{130}Te , the observation of which would demonstrate that the neutrino is a Majorana particle. This measurement is expected to begin in 2017. I have presented in detail the methods by which the sensitivity of this search is estimated, based on a counting analysis analogous to the SNO *hep/DSNB* search. Beyond this, I have presented an overview of relevant systematic uncertainties based on a series of sideband analyses developed to provide robust data-driven constraints on the dominant uncertainties.

The expected sensitivity for the initial phase of SNO+ is $T_{1/2}^{0\nu} > 9.4 \times 10^{25}$ y at the 90% confidence level, assuming a five-year run. This corresponds to a limit on the effective Majorana neutrino mass $m_{\beta\beta}$ in the range 55.2 – 133 meV, depending on the choice of matrix element. This range is competitive with current-generation experiments, and covers the parameter range near the top of the inverted neutrino mass hierarchy.

Following this Phase I run, the detector will be upgraded with a high-quantum efficiency PMT array and the Te loading increased tenfold to 3.0%. This leads to nearly an order of magnitude in half-life sensitivity, pushing the 90% CL limit to $T_{1/2}^{0\nu} > 8.4 \times 10^{26}$ y (also for five live-years), or $m_{\beta\beta}$ in the range 19.0 – 45.9 meV. This unprecedented sensitivity would cover the vast majority of the inverted mass hierarchy parameter space.

The sensitivity of SNO+ Phases I and II are illustrated in Figure A.6, along with selected results from GERDA and KamLAND-Zen as well as the outstanding claim of observation in ^{76}Ge by Klapdor-Kleingrothaus and Krivosheina [125].

SNO+ will be the first experiment to load Te into liquid scintillator, and beyond making a competitive measurement in Phase I and achieving groundbreaking sensitivity in Phase II, will serve to demonstrate the power of this approach for future, larger-scale experiments using liquid scintillator and water-based liquid scintillator (WbLS) targets.

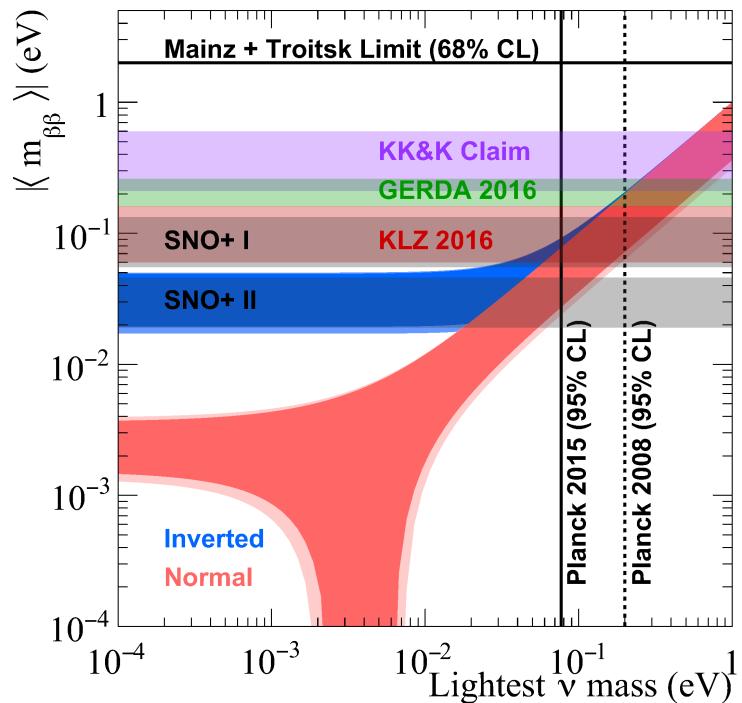


Figure A.6: A Vissani-Strumia plot showing the parameter space for $0\nu\beta\beta$ in the canonical three-neutrino mixing model, with the projected sensitivity for SNO+ as well as selected experimental limits (see Section A.1.1 for details).

Appendix B

Atmospheric Neutrino Event

Simulation with GENIE

The GENIE (Generates Events for Neutrino Interaction Experiments) generator [99] is a software package which simulates high-energy (~ 100 MeV to 100 GeV) neutrino interactions with nuclear targets in a geometric detector model, using state-of-the-art modeling. We may input to GENIE a calculated atmospheric neutrino flux and a simplified model of the SNO detector to obtain a sample of primary particle tracks, and then import these tracks into the full detector simulation (SNOMAN) to generate, propagate, and detect optical photons with the appropriate response. GENIE also includes utilities for reweighting the simulation output for variations in relevant model parameters, and important feature for evaluating model systematics.

B.1 Physics Models

The nuclear physics model used for all processes is the Bodek and Ritchie relativistic Fermi gas model which accounts for short-range correlations and has been verified by electron scattering measurements [153], with an energy- and A -dependent nuclear density. For the quasi-elastic scattering cross section, the Llewellyn-Smith model [154] is used, with the BBBA2005 parameterization of the electromagnetic form factors [155]. The elastic neutral current scattering and baryon resonance production cross sections are due to Arhens et al. [156] and Rein-Sehgal [157], respectively. Non-resonance inelastic scattering, including deep inelastic scattering (DIS) is implemented with the Bodek and Yang model [158]. Further details on the cross section modeling may be found in [99].

The production of hadrons in GENIE is done with the specially-developed AGKY model, which smoothly transitions from a Koba-Nielsen-Olesen [159] model at low invariant mass to the widely-used PYTHIA/JETSET [160] model at high mass. These hadrons may rescatter within the nucleus; GENIE relies on the INTRANUKE subpackage for simulating intranuclear pion and nucleon transport and final state interactions (FSI), described in [99].

B.2 Detector Model

A simplified SNO geometry was created in ROOT for use in the GENIE simulation. This model includes:

Acrylic Vessel The AV is modeled as a union of a spherical shell (5 cm thick and having a radius of 600 cm) and a cylindrical shell. The material is assumed to be $C_5H_8O_2$.

Light Water The AV is contained within an 8.5 m sphere of H_2O . The PMTs and support structure are not included.

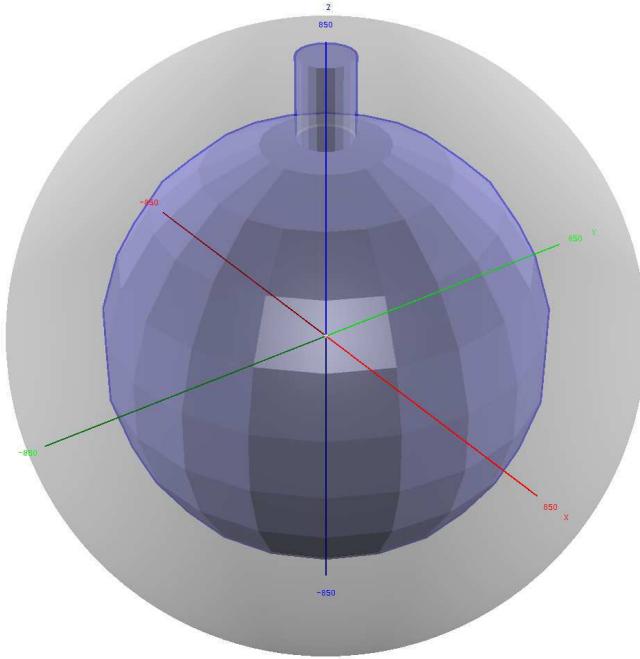


Figure B.1: Visualization of the simplified SNO geometry used in the GENIE primary vertex generation. The detector model includes the acrylic vessel with neck (blue), light water (gray), and heavy water. The facets are an artifact of the visualization.

Heavy Water Inside the vessel is an interchangeable target, which can be pure D_2O , D_2O loaded with 0.2% natural $NaCl$, or D_2O with an array of NCDs, depending on the phase.

A visualization of the geometry is shown in Figure B.1.

Appendix C

A Time-Correlated Approach to Supernova Background Detection

The general discussion of the DSNB presented in this dissertation has assumed, naturally, that it is diffuse: isotropic and uniformly distributed in space and time. However, for a small fraction of the volume nearest to Earth, with $z \lesssim 0.05$, a significant number of Type II supernovae occurring during SNO were observed optically, and so the time of the event is approximately known, in addition to (in most cases) the direction and distance to the host galaxy. In terms of flux, this volume corresponds to about 10% of the total DSNB signal from 20–30 MeV, according to the Beacom and Strigari model [63, 65]. By extrapolating the light curve back to estimate the probable arrival time for neutrinos, it is possible to perform a time-correlated search with significantly reduced background.

To determine the practicality of such an approach, let us perform a rough estimate. According to the model in Reference [161], the core-collapse supernova rate in this low- z

Reaction	Region	Expected Counts		
		Phase I	Phase II	Phase III
CC $\bar{\nu}_e + p \rightarrow n + e^+$	H ₂ O	409.71	409.71	409.71
CC $\bar{\nu}_e + d \rightarrow n + n + e^+$	D ₂ O	135.88	135.88	135.88
CC $\nu_e + d \rightarrow p + p + e^-$	D ₂ O	77.74	77.74	77.74
NC $\nu_{\mu\tau} + d \rightarrow \nu_{\mu\tau} + p + n$	D ₂ O	83.14	188.54	120.78
NC $\bar{\nu}_{\mu\tau} + d \rightarrow \bar{\nu}_{\mu\tau} + p + n$	D ₂ O	67.34	152.26	97.40

Table C.1: Number of time-integrated neutrino events in each SNO phase for dominant channels, above a threshold of 16 hit PMTs, according to the Beacom and Vogel model [165]. Data from Reference [164], Table 7.1.

region is approximately flat, $\sim 1.5 \times 10^{-4} \text{ y}^{-1} \text{ Mpc}^{-3}$. If we efficiently sample SN with $z \lesssim 0.03$, this implies about 600 such events are expected during SNO. In the supernova catalog provided by the IAU Central Bureau for Astronomical Telegrams [162], 471 SNII are identified, 311 of which are associated with a known astronomical object with a measured redshift recorded in the NASA/IPAC Extragalactic Database (NED)²³ [163].

Overall, the dominant channel for supernova neutrino interactions in SNO is $\bar{\nu}_e$ CC on protons in the external H₂O region ($\bar{\nu}_e + p \rightarrow n + e^+$), followed by $\nu_{\mu\tau} - d$ and $\bar{\nu}_{\mu\tau} - d$ NC in the D₂O, then $\nu_e - d$ and $\bar{\nu}_e - d$ CC [164]. The expected event rates for a supernova at a distance of 10 kpc are given in Table C.1. Relative to the 10 kpc benchmark, the number of events detected scales with the flux, as $1/r^2$ where r is the distance to the supernova. Using redshift data from the NED database for SNII that occurred while SNO was taking data, the number of events for each channel C is given by

$$N_C^{(10 \text{ kpc})} \sum_i^{N_{\text{SN}}} \eta_i \left(\frac{10 \text{ kpc}}{z_i c / H_0} \right)^2 \quad (\text{C.1})$$

²³The NASA/IPAC Extragalactic Database (NED) is operated by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

where η is a detection efficiency including analysis cuts.

T. Shokair has previously performed a comprehensive analysis of the full SNO dataset to search generically for antineutrinos and identified 47 events with a background of 42.68 ± 5.02 [166]. Before any analysis cuts, about 0.14 events due to $\bar{\nu}$ are expected, much smaller than the systematic uncertainty in the background estimate.

In the CC $\nu_e - d$ channel, about 2.4×10^{-3} events are expected in the range 20–30 MeV for the Beacom and Vogel model [165], with an in-time background of 0.24 events, assuming a one-day search window for each candidate SN and a 5.5 m fiducial volume. As expected from the Beacom and Strigari DSNB model, these “nearby” events contribute about 10% of the expected DSNB flux²⁴ (0.021 events). The signal-to-background ratio is about 10^{-2} , and the background-fluctuation sensitivity ~ 900 times the signal expectation, which is not improved significantly by further tightening the search window as it is already driven by low signal statistics. As this is much weaker than the sensitivity achievable in the averaged approach where the statistics are somewhat better, we do not pursue the cross-correlation of SNO data with the SN catalog further. However, this approach may still be of interest for more sensitive detectors.

²⁴Note that the averaged DSNB approach does include the flux due to the nearby SN, but averaged over time and space.

Appendix D

Low-Level Cuts

The “low-level” cuts are applied upstream of any reconstruction or analysis, and are primarily intended to remove background events caused by detector effects (e.g. flashing PMTs, electronic pickup). They also tag events based on time, for example events closely following a tagged muon.

Retrigger Event occurred less than $5 \mu\text{s}$ after the previous event

Charge vs. Time (QvT) Highest-charge PMT has very high charge and occurs early.

Charge vs. N_{hit} (Q/NHIT) Too large a fraction of the total charge of the event occurred in too few PMTs.

Crate Isotropy Hit electronics channels are too localized: 70% within a crate and 80% of those on adjacent FECs.

Analog Measurement Board (AMB) The pulse shape characteristics of the ESUMHI trigger pulse are abnormal.

Fitterless Time Spread (FTS) The median time difference of nearby PMTs (which should have similar times) are too large.

Outward-looking PMTs (OWL) More than three outward-looking (OWL) or underwater test tubes (BUTTS) fired in the event.

Junk Event (JUNK) The event fails basic data quality checks, including missing trigger information or duplicate PMT hit data.

Neck Event (NECK) PMTs located in the neck fired during this event (two, or one with high charge).

ESUM Trigger (ESUM) Both ESUMLO and ESUMHI triggers fired, and none of the others generally used for physics triggering.

Charge Cluster (QCLUSTER) A set of nearby PMTs had anomalously high or low charge.

Muon Follower (short) The event came within 20 seconds after a muon event.

In-time Channel Fewer than 60% of the hits fall inside any 93 ns coincidence window.

Flasher Geometry A cluster of hits occurs in either electronics space or PMT space, which is separated from the rest of the hits. This tags “flashers,” which occur when a PMT emits a flash of light (and it and its neighbors see high charge due to pickup) and the light is detected by PMTs across the detector.

OWL Trigger The OWLESUMHI trigger fired.

Missed muon follower (short) The event occurred within 250 ms of another event with $N_{hit} > 60$ (D₂O phase) or $N_{hit} > 150$ (salt/NCD). This cuts followers from events with no visible precursor.

Appendix E

pT Source Run List

The run list for the pT source data is found in SNO internal note MANN-7HJP7R, “A study on the pT source data,” by Feng Zhang. For posterity, the runs are:

(0,0,0) cm 15172–15177, 15235, 15237–15242

(0,0,-250) cm 15263–15266

(0,0,-500) cm 15252–15254, 15257–15259, 15261

Appendix F

^8Li Source Run List

The following ^8Li source runs (with the given source positions) were selected for use in this analysis based on information reported by shift operators.

Phase I

14348 $(-16.0, 23.5, 0.08)$

14371 $(-16.0, 23.5, 249.14)$

14373 $(-16.0, 23.5, 550.04)$

Phase II

23054, 23056, 23058, 23071, 23108 $(0.0, -21.6, 0.0)$

23073, 23075, 23077 $(0.0, -21.6, -254.5)$

23084, 23086, 23088, 23091 $(384.9, -1.2, 73.9)$

23104, 23105 (0.0, -21.6, 249.8)

23094 (0.0, -21.6, 450.1)

In Runs 23073–23077, argon was (accidentally) used instead of nitrogen as a wavelength shifter; PSD is still possible, so these runs are included.

Phase III The Phase III ${}^8\text{Li}$ run list is taken from SNO internal note MANN-7GFKTL, “A study on the Energy Nonlinearity using ${}^8\text{Li}$ source of NCD phase dataset,” by Feng Zhang. Runs with similar source locations are collected into groups.

Group 1 (0, 0, 0) **cm** 63334 63341 63342 63344 63346 63362 63369 63371

Group 2 (0, 0, 200) **cm** 63396 63398

Group 3 (4, 1, 258) **cm** 63401 63403 63405 63407 63415 63417

Group 4 (0, 0, 400) **cm** 63432 63434 63444 63449

Group 5 (-384, 0, 76) **cm** 63381 63382 63383 63388 63394

Group 6 (384, 0, 76) **cm** 63374 63375 63379

Appendix G

Frequentism

One straightforward way to define uncertainty is to repeat an experiment many, many times, and use the deviations in the results to characterize the inherent errors. Assuming that the true value μ of a physical parameter we seek to measure has a fixed but unknown value, we may define a range of values $[\mu_1, \mu_2]$ such that

$$P(\mu \in [\mu_1, \mu_2]) = \alpha \quad (\text{G.1})$$

This is a statement about the behavior of an ensemble of many experiments: we define the interval $[\mu_1, \mu_2]$ such that in a fraction α of experiments, the interval will contain the true value μ .²⁵ Importantly, this is not a statement about the probability distribution of μ itself, which in the frequentist interpretation does not exist. Hence, it is not correct to assume that there is a probability α that μ lies inside any particular range of values.

²⁵It is a common misconception that the experiments forming an ensemble must be “identical” and that this makes the frequentist approach impractical. In fact the only requirement is that all measure μ .

Assuming that one cannot run many experiments in order to define the intervals exactly using data — which is indeed the case for this analysis, given SNO’s currently-unique sensitivity to the *hep* and DSNB ν_e signals — how might one construct an interval given a single measurement?

The Neyman construction [118] provides a general framework, but allows the experimenter a good deal of freedom. For all possible values of μ , one builds a probability distribution for the measured parameter $\hat{\mu}$, and defines an interval in $\hat{\mu}$ that has a desired probability of occurring. Then, one draws a “belt” across these intervals, taking as μ_1 and μ_2 the values of μ that bound the allowed region given an observed $\hat{\mu}$. The choice of an ordering principle to determine which points in each $P(n|\hat{\mu})$ distribution are added to the belt for $\hat{\mu}$ is left to the experimenter. Simply ordering by probability is often referred to as the classical frequentist approach, though the approach is in fact more general.

The Feldman-Cousins “unified” approach is an extension of the classical frequentist method, representing a particular ordering principle for choosing the belts [110]. Points are added in order of the likelihood ratio:

$$R = \frac{P(n; \mu)}{P(n; \tilde{\mu})} \tag{G.2}$$

where $\tilde{\mu}$ is the best-fit allowed value of μ . This choice of ordering naturally transitions between one-sided and two-sided confidence intervals, avoiding the “flip-flopping” under-coverage issues inherent with a choice of interval based on the data, and is an intuitive choice when considered in the context of hypothesis testing (see, e.g., Reference [167]). For these reasons, the Feldman-Cousins approach has become very common in high-energy physics applications, and is adopted for the analysis presented in this dissertation as well, for consistent comparison with previous results and relative ease of interpretation.

References

- [1] BRUCE T CLEVELAND, ET AL. **Measurement of the Solar Electron Neutrino Flux with the Homestake Chlorine Detector.** *The Astrophysical Journal*, **496**(1):505–526, March 1998. 1, 29
- [2] W HAMPEL, ET AL. **GALLEX solar neutrino observations: results for GALLEX IV.** *Physics Letters B*, **447**(1-2):127–133, February 1999. 31
- [3] M ALTMANN, ET AL. **Complete results for five years of GNO solar neutrino observations.** *Phys. Lett. B*, **616**(hep-ex/0504037. MPI-K-2005-008):174–190. 22 p. 2005. 31
- [4] SAGE COLLABORATION AND J N ABDURASHITOV. **Measurement of the Solar Neutrino Capture Rate by the Russian-American Gallium Solar Neutrino Experiment During One Half of the 22-Year Cycle of Solar Activity.** *Journal of Experimental and Theoretical Physics*, **95**(2):12–12, 2002. 31
- [5] THE SUPER-KAMIOKANDE COLLABORATION. **Measurement of the solar neutrino energy spectrum using neutrino-electron scattering.** *Physical Review Letters*, **82**:12–12, 1998. 31
- [6] Q R AHMAD, ET AL. **Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory.** *Physical Review Letters*, **87**(7):071301, August 2001. 5, 21, 31
- [7] G BELLINI, ET AL. **Final results of Borexino Phase-I on low-energy solar neutrino spectroscopy.** *Physical Review D*, **89**(1):112007, June 2014. 1, 215
- [8] F P AN, ET AL. **Observation of Electron-Antineutrino Disappearance at Daya Bay.** *Physical Review Letters*, **108**(17):171803, April 2012. 1
- [9] J K AHN, ET AL. **Observation of Reactor Electron-Antineutrinos Disappearance in the RENO Experiment.** *Physical Review Letters*, **108**(19):191802, May 2012.
- [10] Y ABE, ET AL. **Indication of Reactor $\bar{\nu}_e$ Disappearance in the Double Chooz Experiment.** *Physical Review Letters*, **108**(13):131801, March 2012.
- [11] P ADAMSON, ET AL. **Electron Neutrino and Antineutrino Appearance in the Full MINOS Data Sample.** *Physical Review Letters*, **110**(17):171801, April 2013. 1
- [12] JOHN N BAHCALL. *Neutrino astrophysics.* Cambridge University Press, Institute for Advanced Study, Princeton, NJ, 1989. 1, 31
- [13] W DAVID ARNETT, JOHN N BAHCALL, ROBERT P KIRSHNER, AND STANFORD E WOOSLEY. **Supernova 1987A.** *Annual Review of Astronomy and Astrophysics*, **27**(1):629–700, September 1989. 2, 35
- [14] K BAYS, ET AL. **Supernova relic neutrino search at super-Kamiokande.** *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **85**(5):052007, March 2012. 2, 88, 170
- [15] THE SNO COLLABORATION. **A Search for Neutrinos from the Solar hep Reaction and the Diffuse Supernova Neutrino Background with the Sudbury Neutrino Observatory.** *The Astrophysical Journal*, **653**:12–12, 2006. xvii, 2, 32, 34, 43, 79, 89, 90, 97, 98, 103, 107, 116, 136, 149, 166, 211
- [16] C HOWARD. *A Search for hep Neutrinos with the Sudbury Neutrino Observatory Detector.* PhD thesis, Edmonton, Alberta. 2, 79, 92, 121, 124, 205
- [17] WOLFGANG ERNST FRIEDERICH PAULI. **Open letter to radioactive persons.** 1930. 5
- [18] C L JR COWAN, F REINES, F B HARRISON, H W KRUSE, AND A D MCGUIRE. **Detection of the Free Neutrino: A Confirmation.** *Science*, **124**(3):103–104, July 1956. 5

[19] Y FUKUDA, ET AL. **Evidence for Oscillation of Atmospheric Neutrinos.** *Physical Review Letters*, **81**(8):1562–1567, August 1998. 5, 10, 31, 81

[20] STEVEN WEINBERG. **A Model of Leptons.** *Physical Review Letters*, **19**(21):1264–1266, November 1967. 7

[21] CARLO GIUNTI AND KIM CHUNG WOOK. *Fundamentals of Neutrino Physics and Astrophysics*. Oxford University Press, Oxford, 2007. xiii, xvi, 8, 23, 24, 26, 37, 83, 86

[22] G AAD, ET AL. **Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC.** *Physics Letters B*, **716**(1):1–29, September 2012. 10

[23] S CHATRCHYAN, ET AL. **Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC.** *Physics Letters B*, **716**(1):30–61, September 2012. 10

[24] BORIS KAYSER. **On the quantum mechanics of neutrino oscillation.** *Physical Review D (Particles and Fields)*, **24**(1):110–116, July 1981. 14

[25] L WOLFENSTEIN. **Neutrino oscillations in matter.** *Physical Review D (Particles and Fields)*, **17**(9):2369–2374, May 1978. 14

[26] S P MIKHEYEV AND A YU SMIRNOV. **Resonant amplification of ν oscillations in matter and solar-neutrino spectroscopy.** *Il Nuovo Cimento C*, **9**(1):17–26, 1986. 14

[27] S R SEIBERT. *A Low Energy Measurement of the 8B Solar Neutrino Spectrum at the Sudbury Neutrino Observatory*. PhD thesis, Austin, TX, July 2008. xvii, 15, 48, 62, 106, 129, 174

[28] K NAKAMURA AND PARTICLE DATA GROUP. **Review of Particle Physics.** *Journal of Physics G: Nuclear and Particle Physics*, **37**(7):075021, July 2010. 16, 50, 103, 218

[29] M C GONZALEZ-GARCIA, MICHELE MALTONI, AND THOMAS SCHWETZ. **Updated fit to three neutrino mixing: status of leptonic CP violation.** *Journal of High Energy Physics*, **2014**(11):1–28, November 2014. xiii, 17, 18

[30] K A OLIVE AND OTHERS. **Review of Particle Physics.** *Chin. Phys. C*, **38**(9):090001, 2014. 17

[31] N GREVESSE AND A J SAUVAL. **Standard Solar Composition.** *Space Science Reviews*, **85**(1):161–174, May 1998. 22

[32] E G ADELBERGER, ET AL. **Solar fusion cross sections. II. The pp chain and CNO cycles.** *Review of Modern Physics*, **83**(1):195–246, January 2011. 22

[33] JOHN N BAHCALL, ALDO M SERENELLI, AND SARBANI BASU. **New Solar Opacities, Abundances, Helioseismology, and Neutrino Fluxes.** *The Astrophysical Journal*, **621**(1):L85–L88, March 2005. xvi, 25

[34] JOHN N BAHCALL, M H PINSONNEAULT, AND SARBANI BASU. **Solar Models: Current Epoch and Time Dependences, Neutrinos, and Helioseismological Properties.** *The Astrophysical Journal*, **555**(2):990–1012, July 2001. 24

[35] JOHN N BAHCALL AND M H PINSONNEAULT. **What Do We (Not) Know Theoretically about Solar Neutrino Fluxes?** *Physical Review Letters*, **92**(12):121301, March 2004. 24

[36] JOHN N BAHCALL, ALDO M SERENELLI, AND SARBANI BASU. **10,000 Standard Solar Models: A Monte Carlo Simulation.** *The Astrophysical Journal Supplement Series*, **165**(1):400–431, July 2006. 24

[37] R BONVENTRE, ET AL. **Nonstandard models, solar neutrinos, and large θ_{13} .** *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **88**(5):053010, September 2013. 27

[38] N BARROS. *Precision Measurement of Neutrino Oscillation Parameters*. PhD thesis, Universidade de Lisboa, March 2013. xvi, 28, 64, 65

[39] RAYMOND DAVIS. **Attempt to Detect the Antineutrinos from a Nuclear Reactor by the $\text{Cl}^{37}(\bar{\nu}, e^-)\text{Ar}^{37}$ Reaction.** *Physical Review*, **97**(3):766–769, February 1955. 29

[40] YOJI TOTSUKA. **Recent results on solar neutrinos from Kamiokande.** *Nuclear Physics B*, **19**:69–76, April 1991. 31

[41] SAGE COLLABORATION. **Measurement of the solar neutrino capture rate with gallium metal.** *Physical Review C*, **60**:34–34, 1999. 31

[42] F KAETHER, W HAMPEL, G HEUSSER, J KIKO, AND T KIRSTEN. **Reanalysis of the Gallex solar neutrino flux and source experiments.** *Physics Letters, Section B: Nuclear, Elementary Particle and High-Energy Physics*, **685**(1):47–54, 2010. 31

[43] Q R AHMAD, ET AL. **Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory.** *Physical Review Letters*, **89**(1):011301, June 2002. 31, 91

[44] B AHARMIM, ET AL. **Determination of the ν_e and total B8 solar neutrino fluxes using the Sudbury Neutrino Observatory Phase I data set.** *Physical Review C - Nuclear Physics*, **75**(4):045502, April 2007. 67, 68, 69, 107

[45] B AHARMIM, ET AL. **Electron energy spectra, fluxes, and day-night asymmetries of 8B solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory.** *Physical Review C - Nuclear Physics*, **72**(5):055502, November 2005. 80, 91, 98, 102

[46] B AHARMIM, ET AL. **Low-energy-threshold analysis of the Phase I and Phase II data sets of the Sudbury Neutrino Observatory.** *Physical Review C - Nuclear Physics*, **81**(5):055504, May 2010. 70, 80, 91, 94, 96, 100, 106, 111, 138, 196, 223

[47] B AHARMIM, ET AL. **Measurement of the ν_e and total 8B solar neutrino fluxes with the Sudbury Neutrino Observatory phase-III data set.** *Physical Review C - Nuclear Physics*, **87**(1):015502, January 2013. 67, 68, 106

[48] B AHARMIM, ET AL. **Combined analysis of all three phases of solar neutrino data from the Sudbury Neutrino Observatory.** *Physical Review C - Nuclear Physics*, **88**(2):025501, August 2013. 31, 65, 92, 93, 100, 102, 123, 138, 150, 234

[49] Y FUKUDA, ET AL. **Measurements of the Solar Neutrino Flux from Super-Kamiokande's First 300 Days.** *Physical Review Letters*, **81**(6):1158–1162, 1998. 31

[50] J HOSAKA, ET AL. **Solar neutrino measurements in Super-Kamiokande-I.** *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **73**(11):112001, June 2006. 32

[51] J P CRAVENS, ET AL. **Solar neutrino measurements in Super-Kamiokande-II.** *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **78**(3):1–32, 2008.

[52] K ABE, ET AL. **Solar neutrino results in Super-Kamiokande-III.** *Physical Review D*, **83**(5):052010, March 2011. 31

[53] C ARPESELLA, ET AL. **Direct measurement of the Be7 solar neutrino flux with 192 days of Borexino data.** *Physical Review Letters*, **101**(9):1–6, 2008. 31

[54] G BELLINI, ET AL. **First Evidence of pep Solar Neutrinos by Direct Detection in Borexino.** *Physical Review Letters*, **108**(5):051302, February 2012. 31

[55] BOREXINO COLLABORATION, ET AL. **Neutrinos from the primary proton-proton fusion process in the Sun.** *Nature*, **512**(7):383–386, August 2014. 31

[56] S ANDRINGA, ET AL. **Current Status and Future Prospects of the SNO+ Experiment.** *Advances in High Energy Physics*, **2016**, January 2016. 31, 223, 227, 228

[57] J R ALONSO, ET AL. **Advanced Scintillator Detector Concept (ASDC): A Concept Paper on the Physics Potential of Water-Based Liquid Scintillator.** September 2014. 31, 44

[58] Y KOSHIO. **Recent results from Super-Kamiokande on solar neutrinos.** In *Les Rencontres de Physique de la Vallee d'Aoste La Thuile*, Aosta Valley. 32

[59] YOICHIRO SUZUKI. **Solar Neutrinos.** *Lepton and Photon Interactions at High Energies*, page 201, 1999. xvi, 33

[60] S FUKUDA, ET AL. **Constraints on Neutrino Oscillations Using 1258 Days of Super-Kamiokande Solar Neutrino Data.** *Physical Review Letters*, **86**(25):5656–5660, June 2001. xvi, 32, 33

[61] JOHN F BEACOM. **The Diffuse Supernova Neutrino Background.** *Annual Review of Nuclear and Particle Science*, **60**:25–25, 2010. 37, 39

[62] MATHIAS TH KEIL, GEORG G RAFFELT, AND HANS-THOMAS JANKA. **Monte Carlo Study of Supernova Neutrino Spectra Formation.** *The Astrophysical Journal*, **590**(2):971–991, June 2003. 38

[63] JOHN F BEACOM AND LOUIS E STRIGARI. **New test of supernova electron neutrino emission using Sudbury Neutrino Observatory sensitivity to the diffuse supernova neutrino background.** *Physical Review C*, **73**(3):035807, March 2006. xvii, 38, 41, 42, 211, 241

[64] PLANCK COLLABORATION. **Planck 2015 Results. XIII. Cosmological parameters.** *Astronomy & Astrophysics*, pages 20–20, 2015. xiii, 38, 39, 218, 222

[65] LOUIS E STRIGARI, JOHN F BEACOM, TERRY P WALKER, AND PENGJIE ZHANG. **The concordance cosmic star formation rate: implications from and for the supernova neutrino and gamma ray backgrounds.** *Journal of Cosmology and Astroparticle Physics*, **2005**(04):017–017, April 2005. xvii, 38, 40, 241

[66] TOMAS DAHLEN, ET AL. **High-Redshift Supernova Rates.** *The Astrophysical Journal*, **613**(1):189–199, September 2004. xvii, 40

[67] D SCHIMINOVICH, ET AL. **The GALEX - VVDS Measurement of the Evolution of the Far-Ultraviolet Luminosity Density and the Cosmic Star Formation Rate.** *The Astrophysical Journal*, **619**(1):L47–L50, 2005. xvii, 40

[68] SHAUN COLE, ET AL. **The 2dF galaxy redshift survey: near-infrared galaxy luminosity functions.** *Monthly Notices of the Royal Astronomical Society*, **326**(1):255–273, September 2001. xvii, 40

[69] M AGLIETTA, ET AL. **Limits on low-energy neutrino fluxes with the Mont Blanc liquid scintillator detector.** *Astroparticle Physics*, **1**(1):1–9, November 1992. 43

[70] CECILIA LUNARDINI. **The Diffuse neutrino flux from supernovae: Upper limit on the electron neutrino component from the non-observation of antineutrinos at superkamiokande.** *Phys. Rev.*, **D73**(8):083009, 2006. 44

[71] CECILIA LUNARDINI AND ORLANDO L G PERES. **Upper limits on the diffuse supernova neutrino flux from the SuperKamiokande data.** *Journal of Cosmology and Astroparticle Physics*, **2008**(08):033, August 2008. 44, 170

[72] JOHN F BEACOM AND MARK R VAGINS. **Antineutrino Spectroscopy with Large Water Čerenkov Detectors.** *Physical Review Letters*, **93**(1):171101, October 2004. 44

[73] LIZ KRUESI. **Underground plans.** *Symmetry*, (July 2015). 44

[74] DUNE COLLABORATION, ET AL. **Long-Baseline Neutrino Facility (LBNF) and Deep Underground Neutrino Experiment (DUNE) Conceptual Design Report Volume 2: The Physics Program for DUNE at LBNF.** December 2015. 44

[75] HERBERT H CHEN. **Direct Approach to Resolve the Solar-Neutrino Problem.** *Physical Review Letters*, **55**(14):1534–1536, September 1985. 46

[76] M KAPLINGHAT, G STEIGMAN, AND T P WALKER. **Supernova relic neutrino background.** *Physical Review D*, **62**(4):043001, July 2000. 50

[77] J BOGER, ET AL. **The Sudbury Neutrino Observatory.** *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **449**:172–207, 2000. xvii, 51, 52, 55, 60

[78] JOSHUA R KLEIN, M S NEUBAUER, R VAN BERG, AND F M NEWCOMER. **The SNO Trigger System.** Technical Report SNO-STR-97-035, University of Pennsylvania. xvii, 58

[79] RICHARD JAMES FORD. *Calibration of SNO for the detection of (8)B neutrinos*. PhD thesis, Queen's U., Kingston, Kingston, Ontario, 1999. 61

[80] M R DRAGOWSKY, ET AL. **The 16N calibration source for the Sudbury Neutrino Observatory**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **481**(1-3):284–296, April 2002. 61

[81] N J TAGG, ET AL. **The 8Li calibration source for the Sudbury Neutrino Observatory**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **489**(1-3):178–188, August 2002. 61, 131

[82] M C BROWNE, ET AL. **A Compact ${}^3\text{H}(\text{p},\gamma){}^4\text{He}$ 19.8-MeV Gamma-Ray Source for Energy Calibration at the Sudbury Neutrino Observatory**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **452**(1-2):115–129, 2000. 61

[83] GABRIEL D OREBI GANN. *An Improved Measurement of the 8 B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory*. PhD thesis, Merton College, Oxford, 2008. 62, 100

[84] S J BRICE, ET AL. **First Generation Acrylic Encapsulated U/Th Sources**. Technical Report SNO-STR-99-023, February 2000. 62

[85] W T WINTER, S J FREEDMAN, K E REHM, AND J P SCHIFFER. **The B8 neutrino spectrum**. *Physical Review C - Nuclear Physics*, **73**(2), 2006. 63, 103, 234

[86] WALTER R NELSON, HIDEO HIRAYAMA, AND DAVID W O ROGERS. **The EGS4 Code System**. Technical Report 265, SLAC. 63

[87] MCNP: A General Monte Carlo N-Particle Transport Code. Technical Report LA-12625-M, 1993. 63

[88] **Physics at HERA**. Technical report. 64

[89] P A AARNIO, ET AL. **FLUKA: hadronic benchmarks and applications**. In P DRAGOVITSCH, S L LINN, AND M BURBANK, editors, *MC93 International Conference on Monte Carlo Simulation in High Energy and Nuclear Physics*, pages 88–89, Tallahassee, Florida. 64

[90] MALCOLM BUTLER, JIUNN-WEI CHEN, AND XINWEI KONG. **Neutrino-deuteron scattering in effective field theory at next-to-next-to-leading order**. *Physical Review C - Nuclear Physics*, **63**(3):035501, February 2001. 65

[91] JOHN N BAHCALL, MARC KAMIONKOWSKI, AND ALBERTO SIRLIN. **Solar neutrinos: Radiative corrections in neutrino-electron scattering experiments**. *Physical Review D*, **51**(11):6146–6158, June 1995. 65

[92] J R KLEIN AND M S NEUBAUER. **Using Time and Angle Information in the SNOMAN Path Fitter**. Technical Report SNO-STR-2000-022, June 2016. 67

[93] RYAN FRANCIS MACLELLAN. *The energy calibration for the solar neutrino analysis of all three phases of the Sudbury Neutrino Observatory*. PhD thesis, September 2009. 67, 69, 70

[94] A BIALEK, ET AL. **A rope-net support system for the liquid scintillator detector for the SNO+ experiment**. *Nuclear Inst. and Methods in Physics Research*, **827**:152–160, August 2016. 72

[95] R ALVES, ET AL. **The calibration system for the photomultiplier array of the SNO+ experiment**. *Journal of Instrumentation*, **10**(0):P03002–P03002, March 2015. 72

[96] G D BARR, T K GAISSER, P LIPARI, S ROBBINS, AND T STANEV. **Three-dimensional calculation of atmospheric neutrinos**. *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **70**(2):023006, July 2004. xviii, 84

[97] G BATTISTONI, A FERRARI, T MONTARULI, AND P R SALA. **The atmospheric neutrino flux below 100 MeV: The FLUKA results**. *Astroparticle Physics*, **23**(5):526–534, June 2005. 85

[98] T K GAISSER AND TODOR STANEV. **Path length distributions of atmospheric neutrinos**. *Physical Review D*, **57**(3):1977–1982, February 1998. 86

[99] C ANDREOPoulos, ET AL. **The GENIE neutrino Monte Carlo generator**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **614**(1):87–104, February 2010. 86, 90, 238, 239

[100] H EJIRI. Nuclear deexcitations of nucleon holes associated with nucleon decays in nuclei. *Physical Review C*, **48**(3):1442–1444, September 1993. 88

[101] TAKAAKI KAJITA. **Atmospheric neutrinos**. *New Journal of Physics*, **6**:194–194, 2004. 89

[102] G P ZELLER. **Low Energy Neutrino Cross Sections**. In *AIP Conference Proceedings*, 2004. 89

[103] THE HEP ANALYSIS GROUP. **The hep Analysis: Unified Document**. Technical report. xviii, 90, 104

[104] D CASPER. **The nuance neutrino physics simulation, and the future**. *Nuclear Physics B*, **112**(1-3):161–170, 2002. 90

[105] J R WILSON. *A Measurement of the 8B Solar Neutrino Energy Spectrum at the Sudbury Neutrino Observatory*. PhD thesis, University of Oxford. 91, 95

[106] JOSHUA R KLEIN AND AARON ROODMAN. **Blind Analysis in Nuclear and Particle Physics**. *Annual Review of Nuclear and Particle Science*, **55**(1):141–163, 2005. 99

[107] N TAGG. *The 8Li Calibration Source and Through-Going Muon Analysis in the Sudbury Neutrino Observatory*. PhD thesis, Guelph, Ontario, January 2001. 131

[108] FREDERICK JAMES. *Statistical methods in experimental physics*. World Scientific Publishing Co. Pte. Ltd., Hackensack, NJ, second edition, 2006. 149, 180, 185, 207

[109] RENE BRUN AND FONS RADEMAKERS. **ROOT — An object oriented data analysis framework**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **389**(1-2):81–86, April 1997. 149

[110] GARY J FELDMAN AND ROBERT D COUSINS. **Unified approach to the classical statistical analysis of small signals**. *Physical Review D*, **57**(7):3873–3889, 1998. 154, 160, 224, 250

[111] A A MARKOV. **The theory of algorithms**. *American Mathematical Society Translations*, **15**(2):1–14, 1960. 177

[112] NICHOLAS METROPOLIS, ARIANNA W ROSENBLUTH, MARSHALL N ROSENBLUTH, AUGUSTA H TELLER, AND EDWARD TELLER. **Equation of State Calculations by Fast Computing Machines**. *The Journal of Chemical Physics*, **21**(6):1087–1092, June 1953. 179

[113] W K HASTINGS. **Monte Carlo sampling methods using Markov chains and their applications**. *Biometrika*, **57**(1):97–109, April 1970. 179

[114] ANDREW GELMAN AND DONALD B RUBIN. **Inference from Iterative Simulation Using Multiple Sequences**. *Statistical Science*, **7**(4):457–472, November 1992. 181

[115] A E RAFTERY AND S LEWIS. **How many iterations in the Gibbs sampler**. *Bayesian statistics*, 1992. 181

[116] MARY KATHRYN COWLES AND BRADLEY P CARLIN. **Markov Chain Monte Carlo Convergence Diagnostics: A Comparative Review**. *Journal of the American Statistical Association*, **91**(434):883–904, February 2012. 181

[117] ROBERT D COUSINS. **Why isn't every physicist a Bayesian?** *American Journal of Physics*, **63**(5):398–410, May 1995. 183

[118] J NEYMAN. **Outline of a Theory of Statistical Estimation Based on the Classical Theory of Probability**. *Philosophical Transactions of the Royal Society of London. Series A*, **236**(7):333–380, August 1937. 184, 250

[119] A MASTBAUM. **mastbaum/sxmc** [online]. Available from: <https://github.com/mastbaum/sxmc>. 189

[120] NVIDIA CORPORATION. **CUDA C Programming Guide** [online]. Available from: http://docs.nvidia.com/cuda/pdf/CUDA_C_Programming_Guide.pdf. 190

[121] MARK HARRIS. **hemi** [online]. Available from: <https://github.com/harrism/hemi>. 191

[122] HEIKKI HAARIO, EERO SAKSMAN, AND JOHANNA TAMMINEN. **An Adaptive Metropolis Algorithm**. *Bernoulli*, **7**(2):223–242, 2001. 192

[123] M YEH, ET AL. **A new water-based liquid scintillator and potential applications**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **660**(1):51–56, 2011. 215, 228

[124] J SCHECHTER AND J W F VALLE. **Neutrinoless double- β decay in $SU(2) \times U(1)$ theories**. *Physical Review D*, **25**(11):2951, June 1982. 216

[125] H V Klapdor-Kleingrothaus AND I V Krivosheina. **THE EVIDENCE FOR THE OBSERVATION OF $0\nu\beta\beta$ DECAY: THE IDENTIFICATION OF $0\nu\beta\beta$ EVENTS FROM THE FULL SPECTRA**. *Modern Physics Letters A*, **21**(20):1547–1566, 2006. 218, 220, 236

[126] R G H ROBERTSON. **EMPIRICAL SURVEY OF NEUTRINOLESS DOUBLE BETA DECAY MATRIX ELEMENTS**. *Modern Physics Letters A*, **28**(08):1350021–1350021, 2013. 220

[127] C E AALSETH, ET AL. **COMMENT ON "EVIDENCE FOR NEUTRINOLESS DOUBLE BETA DECAY"**. *Modern Physics Letters A*, **17**(22):1475–1478, July 2002. 220

[128] M AGOSTINI, ET AL. **Results on Neutrinoless Double- β Decay of ^{76}Ge from Phase I of the GERDA Experiment**. *Physical Review Letters*, **111**(12):122503, September 2013. 220, 224

[129] H V Klapdor-Kleingrothaus AND I V Krivosheina. **Why is the conclusion of the Gerda experiment not justified**. *Physics of Particles and Nuclei Letters*, **10**(7):704–709, January 2014. 220

[130] R ARNOLD, ET AL. **Measurement of the $\beta\beta$ Decay Half-Life of Te^{130} with the NEMO-3 Detector**. *Physical Review Letters*, **107**(6):062504, August 2011. 221, 222, 233

[131] J B ALBERT, ET AL. **Search for Majorana neutrinos with the first two years of EXO-200 data**. *Nature*, **510**(7504):229–234, June 2014. 221

[132] N ACKERMAN, ET AL. **Observation of Two-Neutrino Double-Beta Decay in Xe^{136} with the EXO-200 Detector**. *Physical Review Letters*, **107**(21):212501, November 2011. 221

[133] J B ALBERT, ET AL. **Improved measurement of the $2\nu\beta\beta$ half-life of ^{136}Xe with the EXO-200 detector**. *Physical Review C*, **89**(1):015502, January 2014. 221

[134] A GANDO, ET AL. **Limit on Neutrinoless $\beta\beta$ Decay of Xe^{136} from the First Phase of KamLAND-Zen and Comparison with the Positive Claim in Ge76**. *Physical Review Letters*, **110**(6):062502, February 2013. 221

[135] KAMLAND-ZEN COLLABORATION. **Search for Majorana Neutrinos near the Inverted Mass Hierarchy region with KamLAND-Zen**. *arXiv.org*, page arXiv:1605.02889, May 2016. 221

[136] F ALESSANDRIA, ET AL. **Sensitivity and Discovery Potential of CUORE to Neutrinoless Double-Beta Decay**. *arXiv.org*, page arXiv:1109.0494, September 2011. 221, 224

[137] K ALFONSO, ET AL. **Search for Neutrinoless Double-Beta Decay of Te^{130} with CUORE-0**. *Physical Review Letters*, **115**(10):102502, September 2015. 221, 224

[138] MAJORANA COLLABORATION, ET AL. **The Majorana Demonstrator: A Search for Neutrinoless Double-beta Decay of ^{76}Ge** . *Journal of Physics: Conference Series*, page 012004, January 2015. 222

[139] S AGOSTINELLI, ET AL. **Geant4—a simulation toolkit**. *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **506**(3):250–303, July 2003. 222

[140] J ALLISON, ET AL. **Geant4 developments and applications**. *IEEE Transactions on Nuclear Science*, **53**(1):270–278, February 2006. 222

[141] S R SEIBERT, ET AL. **RAT-PAC** [online]. Available from: <https://github.com/rat-pac/rat-pac>. 222

[142] C ARPESELLA, ET AL. **Measurements of extremely low radioactivity levels in BOREXINO**. *Astroparticle Physics*, **18**(1):1–25, August 2002. 222

[143] V LOZZA AND J PETZOLDT. **Cosmogenic activation of a natural tellurium target**. *Astroparticle Physics*, **61**:62–71, February 2015. 223

[144] G T EWAN, ET AL. **Sudbury Neutrino Observatory Proposal**. 1987. 223

[145] IAN T COULTER. *Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation*. PhD thesis, Balliol College, Oxford, January 2014. 227

[146] J BAREA, J KOTILA, AND F IACHELLO. **Nuclear matrix elements for double- β decay**. *Physical Review C - Nuclear Physics*, **87**(1), 2013. 229, 231

[147] J KOTILA AND F IACHELLO. **Phase-space factors for double- β decay**. *Physical Review C - Nuclear Physics*, **85**(3):034316, March 2012. 229

[148] FEDOR ŠIMKOVIC, VADIM RODIN, AMAND FAESSLER, AND PETR VOGEL. **$0\nu\beta\beta$ and $2\nu\beta\beta$ nuclear matrix elements, quasiparticle random-phase approximation, and isospin symmetry restoration**. *Physical Review C - Nuclear Physics*, **87**(4):045501, April 2013. 231

[149] J MENENDEZ, A POVES, E CAURIER, AND F NOWACKI. **Disassembling the nuclear matrix elements of the neutrinoless $\beta\beta$ decay**. *Nuclear Physics A*, **818**(3-4):139–151, March 2009. 231

[150] JUHANI HYVÄRINEN AND JOUNI SUHONEN. **Nuclear matrix elements for $0\nu\beta\beta$ decays with light or heavy Majorana-neutrino exchange**. *Physical Review C*, **91**(2):024613, February 2015. 231

[151] TOMÁS R RODRÍGUEZ AND GABRIEL MARTÍNEZ-PINEDO. **Energy Density Functional Study of Nuclear Matrix Elements for Neutrinoless $\beta\beta$ Decay**. *Physical Review Letters*, **105**(25):252503, December 2010. 231

[152] MATTHEW REDSHAW, BRIANNA J MOUNT, EDMUND G MYERS, AND FRANK T III AVIGNONE. **Masses of Te130 and Xe130 and Double- β -Decay Q Value of Te130**. *Physical Review Letters*, **102**(2):212502, May 2009. 234

[153] A BODEK AND J L RITCHIE. **Further studies of Fermi-motion effects in lepton scattering from nuclear targets**. *Physical Review D*, **24**(5):1400–1402, September 1981. 239

[154] C H LLEWELLYN SMITH. **Neutrino reactions at accelerator energies**. *Physics Reports*, **3**(5):261–379, June 1972. 239

[155] R BRADFORD, A BODEK, H BUDD, AND J ARRINGTON. **A New Parameterization of the Nucleon Elastic Form Factors**. *Nuclear Physics B*, **159**:127–132, September 2006. 239

[156] L A AHRENS, ET AL. **Measurement of neutrino-proton and antineutrino-proton elastic scattering**. *Physical Review D*, **35**(3):785–809, February 1987. 239

[157] DIETER REIN AND LALIT M SEHGAL. **Neutrino-excitation of baryon resonances and single pion production**. *Annals of Physics*, **133**(1):79–153, April 1981. 239

[158] A BODEK AND U K YANG. **Higher twist, ξ_W scaling, and effective LO PDFs for lepton scattering in the few GeV region**. *Journal of Physics G: Nuclear and Particle Physics*, **29**(8):1899–1905, August 2003. 239

[159] Z KOBA, H B NIELSEN, AND P OLESEN. **Scaling of multiplicity distributions in high energy hadron collisions**. *Nuclear Physics B*, **40**:317–334, 1972. 239

[160] TORBJÖRN SJÖSTRAND, STEPHEN MRENNNA, AND PETER SKANDS. **A brief introduction to PYTHIA 8.1**. *Computer Physics Communications*, **178**(11):852–867, June 2008. 239

[161] SHUNSAKU HORIUCHI, JOHN F BEACOM, AND ELI DWEK. **Diffuse supernova neutrino background is detectable in Super-Kamiokande.** *Physical Review D - Particles, Fields, Gravitation and Cosmology*, **79**(8):1–14, 2009. 241

[162] **List of Supernovae** [online]. Available from: <http://www.cbat.eps.harvard.edu/lists/Supernovae.html>. 242

[163] **NASA/IPAC Extragalactic Database - NED** [online]. Available from: <https://ned.ipac.caltech.edu/>. 242

[164] J HEISE. *A Search for Supernova Neutrinos with the Sudbury Neutrino Observatory*. PhD thesis, The University of British Columbia, December 2001. xv, 242

[165] J F BEACOM AND P VOGEL. **Mass signature of supernova ν_μ and ν_τ neutrinos in the Sudbury Neutrino Observatory**. *Physical Review D*, **58**(9), 1998. xv, 242, 243

[166] T M SHOKAIR. *TIME-CORRELATED COINCIDENCES AT THE SUDBURY NEUTRINO OBSERVATORY: AN ANTINEUTRINO SEARCH*. PhD thesis, University of Pennsylvania, 2012. 243

[167] ALAN STUART AND J KEITH ORD. *Kendall's advanced theory of statistics. Vol. 2*. The Clarendon Press, Oxford University Press, New York, fifth edition, 1991. 250