Abstract

Measurement of the Reactor Antineutrino Spectrum of ^{235}U by PROSPECT and Daya Bay

Jeremy Kalaheo Kaikane Gaison

2021

Since their experimental discovery over 60 years ago, neutrinos have proven to be a fascinating means of exploring the physical universe. Through a variety of both natural and man-made sources, physicists have discovered many unusual features about these particles from their oscillation between their different flavor states to their particularly small mass. There are still many questions to answer regarding these fundamental particles, though. Among these is whether a possible fourth type of neutrino exists, a sterile neutrino, which could resolve a range of discrepancies between recent measurements and predictions at different energies and baselines. Precision neutrino measurements may also help serve as a benchmark for nuclear physics and shed light on recently discovered spectral anomalies.

The PROSPECT experiment is a 4 ton, ^{6}Li-loaded liquid scintillator detector at Oak Ridge National Laboratory operating $<$10m from the High Flux Isotope Reactor, a research reactor highly enriched in ^{235}U, a dominant isotope in nuclear reactor antineutrino production. The segmented design of the detector allows for unprecedented background rejection at the Earth's surface. Here I will describe the design, construction, data taking, and analysis of PROSPECT towards its measurement of over 50,000 neutrino events and the results from both its sterile neutrino search and its precision measurement of the ^{235}U antineutrino spectrum. Further, I will describe the analysis that combines results from PROSPECT and another reactor antineutrino experiment called Daya Bay. The combined datasets produce the most precise measurement of the ^{235}U antineutrino energy spectrum to date.

Measurement of the Reactor Antineutrino Spectrum of ^{235}U by PROSPECT and Daya Bay

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Jeremy Kalaheo Kaikane Gaison

Dissertation Director: Karsten Heeger

December 2021

Acknowledgments

As I've reflected on all the people who have impacted my life in graduate school, I've struggled to decide how to organize these acknowledgements. Part of me wanted to organize this the standard way particle physicists acknowledge dozens of collaborators: alphabetically. But rather than that, I decided that organizing this in a roughly chronological order helped to group some of the key people from early on in my graduate school career towards the beginning while not condemning anyone towards the end to think that they only barely made the cut. Everyone included here has left an undoubtedly positive imprint on my life, and I love you all deeply and without regret.

First to my mother, I can't convey all of the love and support you've given me my entire life. Only in the last few years have I understood the depths of your sacrifices to ensure that I had the opportunities I did to growing up. I would never have been able to pursue science as a passion without your support, your kindness, and your strength.

To my father, there are so many ways that you've helped me to grow and understand the world. Your encouragement has always spurred me on, and you helped me grow into the independent thinker I am today. Your constantly calm demeanor reminds me to take my time when needed, and to consider the impact I may have on others.

To my comrade and wonderful sister Jess, whether it was letting me tag along with you and your friends on camping trips or just talking in silly voices together, you brought so much joy and wonder to growing up. You were always there for me when I needed you, even when things were just as dark for you, and for that I am eternally grateful. Seeing you and Josh raise your kids gives me so much hope for the future, and I love to see my amazing big sister be just as amazing a mother.

To my advisor Karsten Heeger, thank you for the support and trust you've given me across my time in grad school. I admire the combination of your relentless drive to push scientific research to its limits and your constant efforts to make the scientific community

a better place. You guided my academic growth, and you gave me more responsibility than I ever expected to be able to handle but never gave me more than I could. Thank you to Keith Baker, Dave Moore, and Walter Goldberger for your insights and help in elevating this thesis to a level much higher than it would have been without you. Thank you for your time and your ever-welcoming presences through the years. Also thank you to Pedro Ochoa-Ricoux for being my external reader, for your guidance, and for literally helping me get around Beijing where I was completely lost.

To Tom, I can't thank you enough for taking me on under your wing and for what were probably hundreds of hours of chatting about physics, jobs, biscuits, and Hamilton. To Danielle, thank you for the strength and self-confidence you had even within a large collaboration. You showed me what I could aim to be. Without the two of you, I likely wouldn't have even come to Yale in the first place. And to everyone else from that first year in EAL-100, thank you for making working at Wright so amazing: Antonia, Chris, Estella, Field, J1, Jay, Ke, Kyungeun, Luis, Nate, Nikita, Penny, Walter, and Zach.

To Duncan and Paul, thank you for the privilege of living with you at 101 Bishop. Through scotch and rocks, I wouldn't have changed any of it. To Will, closing GPSCY together that first summer was the start of something great. Thank you for your friendship, and for pushing us to actually learn statistics. And to the rest of Squad: Chris, Danny, Mariel, and Ryan, I'm so glad that I got to do this with you. You kept me sane through problem sets and pandemics, and I never would have made it without you. To Sarah, Nicole, and Lev, you helped make year one Monday nights something truly special at trivia, even though we never actually won anything. Our lives went different ways but I'll always remember pumpkin patches, bike rides, and disco panics with you.

To Mary Lou, you challenged me to think deeply and trust in something beyond myself again. With you my life overflowed with joy and love, and you will always be one of my best friends. I pray that your life is filled with peace, jet-setting adventures, and bountiful love and grace.

To Daniel, Diana, and Stephanie, thank you for your unwavering friendship and countless shared experiences. I'll treasure every moment we spent watching movies, playing Animal Crossing, and drinking boba together. To my DnD crew, thank you Emily, Katrina, Phil, Uri, and our fantastic DM Devon for your enthusiasm and your commitment to The Endless and its crew's journey over the course of 3 laughter-filled years.

There are so many amazing mentors that were a part of the PROSPECT experiment, and I cannot thank them enough for their guidance and support. Thank you to Michelle Dolinski, my undergrad advisor who introduced me to this experiment and helped me find my feet in physics research. Thank you to Nathaniel Bowden and Pieter Mumm, who helped oversee my development as a researcher and mentored me in the job search process. Thank you to David Jaffe for wonderful breakfast time conversations and for your endearing wit. And thank you to Chao Zhang and Xin Ji for your invaluable discussions about data unfolding.

To Ben, my friend, colleague, and confidant on all things joint analysis. It made it so much easier going through this with someone else, and I wish you nothing but smooth sailing as you finish up grad school and move on to bigger and better things. To the rest of the matriculating class of 2016, you have made grad school such a wonderful experience, and I am humbled that I may have been even a tiny factor in any of your decisions to attend. Thank you Connor, Emily, Emma, Kelly, Luna, Sam, and Sumita. You are fantastic friends, you are champions, and you inspire me to strive to be better.

To James W, thank you for your endless camaraderie and commiserating when things were going wrong, and thank you for being one of the founders of Endless (not discount) Sushi. To James N, thank you for teaching me that meetings should be brief or non-existent and for always dancing with me in the halls. To Arina, thank you for your bright and uplifting spirit, which shines through even on dark days.

To Frank, Jeff, Kimberly, Lil (happy birthday!), Paula, Tommy, and Victoria, thank you for everything and for keeping Wright Lab running despite all of the chaos we caused.

I'm sure it wasn't easy, and nothing would have ever been accomplished without you.

To everyone who was part of Science in the News throughout the years, thank you for your hard work engaging in science outreach in the community. Special thanks in particular to Alex, Charles, and Darryl for making our first talk, to Elise for co-directing SitN with me, and to Lizzy for not only keeping the program alive but for expanding it in the midst of a pandemic.

Thank you to all the PROSPECT collaborators I haven't had a chance to mention yet, especially everyone who was on the build team or on-site installation: Adam, Corey, Danielle B, Don, Elisa, Henry, Jeremy L, Jason, Michael, Olga, Pranava, Richard, Tim, Tom W, and Xianyi. You all made this collaboration a wonderful work environment, and helped make PROSPECT a reality. I also want to thank Jianrun for his tireless effort on the joint analysis. You truly drove a lot of this work forward, and I'm lucky to have had you as a counterpart on Daya Bay.

Thank you to all the 2019-2020 Graduate Student Assembly E-board members: CJ, Da'Von, Estella, Ethan, Jacob, Jenette, Jo, Julianne, Kristen, Maria, Meaghan, Megan, Nicolle, Ryan, and our amazing chair Lucy. Your selfless work helped foster so many positive changes in the graduate school, and I loved working with you all. Thank you to my church small group for your encouragement and service. Thank you in particular to David and Feifei for opening up your home for us, and to Daniel, Jonathan, and Kierstin for running everything so smoothly.

To Samantha, thank you for joining our group at Yale (although I know you would have also crushed it at Berkeley) and for being such an amazing advocate for others. To Shelly Lesher, thank you not just for an amazing seminar series, but for helping to make everyone always feel included and valuable. You remind us that being a physicist is about more than just research, and it's just as much about being a person.

To Will Johnson, you were an amazing roommate throughout this final academic year. I learned so much from you about robotics and chess, and I deeply valued our late night dis-

cussions on anything and everything. To Sophia, even though we only overlapped briefly at Wright I can't tell you how much I appreciated your friendship, your courageous and resilient energy, and of course all of your swashbuckling exploits.

And last chronologically, but certainly not least, thanks to all the new members of the group. Working from home for a whole year could have meant an anxiety-filled return to lab, but it has been a refreshingly smooth return to "normal." It has been exciting and engaging to interact with all of you as you start your (post/under)graduate careers. Thank you Annie, Claire, Iris, Jorge, Mike, Sophia, Sophie, Talia, and Yuqi. I know you'll continue to make Wright Lab a place worth working for years to come.

To all my grandparents, who loved me unconditionally even from afar.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 History

Neutrinos have a rich history in experimental particle physics. First hypothesized in 1930 by Wolfgang Pauli [1], neutrinos were meant as a resolution to a problem with beta decay. During that time, beta decay was thought of as the process where a neutron, or a larger atomic nucleus containing a neutron, would decay producing a new nucleus with one fewer neutron and one greater proton as well as an electron

$$(A, Z) \rightarrow (A, Z + 1) + e^{-}. \tag{1.1}$$

Being a two body decay, there is an exact solution for the energy of each product in the rest frame of the initial nucleus. Initial measurements of the outgoing e^{-}, however, found a continuous energy distribution as shown in Figure 1.1.

In order to save the principle of conservation of energy, Pauli proposed a third decay product, allowing each product to have a continuous energy spectrum. In order to explain its negligible impact on the maximum energy of the e^{-}, Pauli proposed that it must be massless (or nearly so). Also to conserve conservation of electric charge, it must

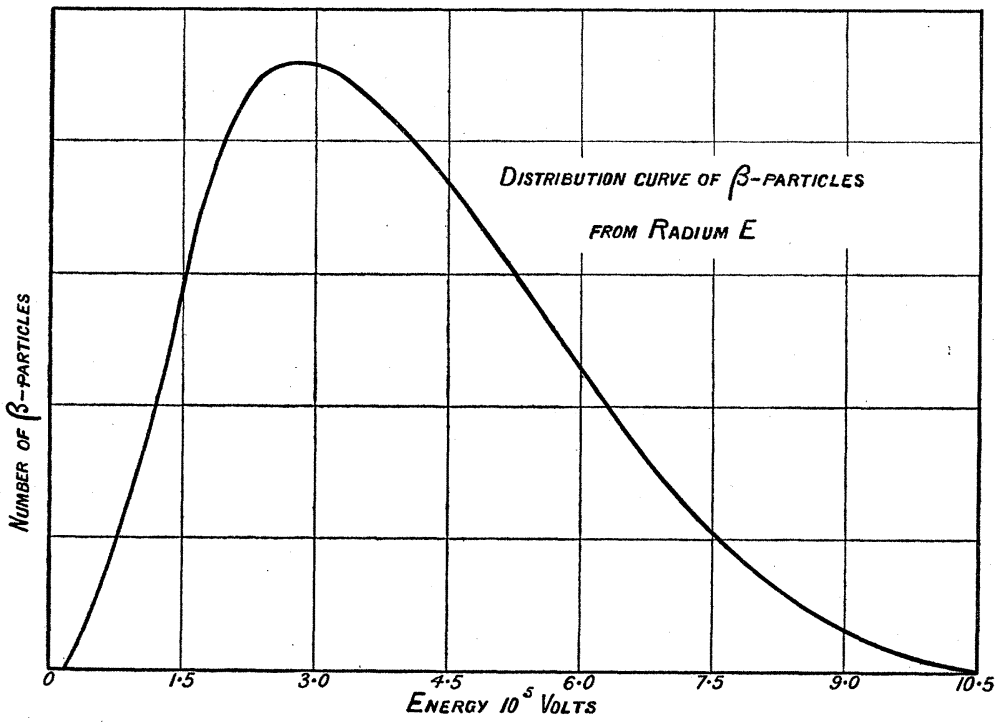

Figure 1.1: Continuous beta decay spectrum of a radium source [2]

be electrically neutral. Pauli's particle, dubbed the neutron at the time but later called the neutrino to distinguish it from the particle discovered by Chadwick [3], was later described by Fermi and predicted to be nearly impossible to detect [4]. Because of the extremely low interaction rate of neutrinos, it was not until after the advent of nuclear reactors and their high antineutrino flux produced by them that they would be experimentally detected. They were first observed in 1956 when Reines and Cowan [5] utilized the process of inverse beta decay (IBD)

$$\overline{\nu}_e + p \rightarrow n + e^+ \tag{1.2}$$

where an incident electron antineutrino interacts with a proton to produce a neutron and a positron, both of which can be measured in coincidence by other more conventional techniques. This experiment was conducted at the Savannah River Plant utilizing a nuclear reactor as a high flux source of antineutrinos. They measured a rate of approximately 3 neutrinos per hour. While consistent with experimental evidence, it wasn't until 1962 and

2000 that there would be direct detections of the muon [6] and tau [7] flavors of neutrinos, respectively, rounding out all of the Standard Model flavors of the neutrino.

1.2 Standard Model Neutrino Interactions

Formally, neutrinos are spin-$\frac{1}{2}$ fermions in the Standard Model of particle physics, a framework of all fundamental particles and their interactions. The three neutrino flavors are coupled with the three flavors of the charged leptons via the weak nuclear force with

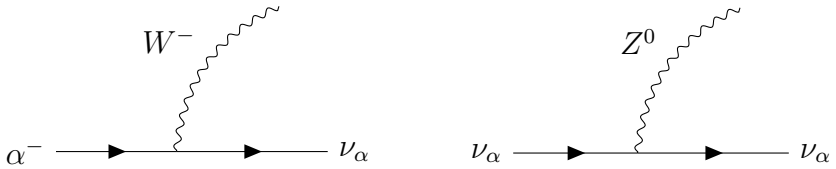

showcasing the charged current and neutral current vertices respectively. Charged current interactions deal with direct interactions between neutrinos and the charged leptons, leading to a vertex like that shown above involving the creation/annihilation of a neutrino and a quark or lepton, and are mediated by the W boson, while neutral current interactions govern scattering interactions involving neutrinos are mediated by the Z boson. Here α corresponds to one of the three lepton generations e, μ, and τ. This leads to one of the primary production mechanism of neutrinos, beta decay, where a neutron decays to a proton p, electron e^-, and an electron antineutrino $\bar{\nu}_e$.

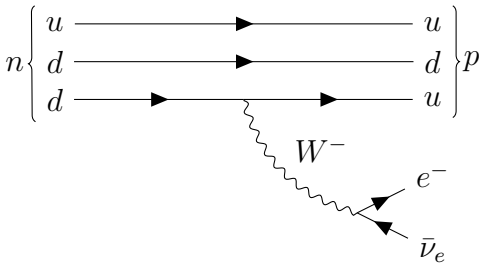

The electron capture process is extremely similar where an electron captures on a proton producing a neutron and an electron neutrino.

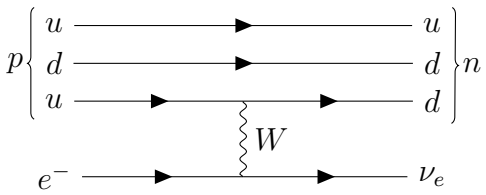

These diagrams can be expressed in condensed versions where the quarks are collapsed into the proton and neutron as

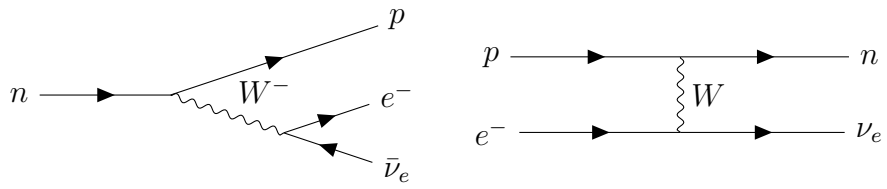

These interactions can occur with protons and neutrons in free space, but may also deal with protons and neutrons in atomic nuclei. If beta decay (electron capture) occurs in a nucleus, the atom will maintain the same total mass number, the number of protons will increase (decrease) while the number of neutrons in the nucleus will decrease (increase). Additionally, kinetic energy will be imparted on the outgoing products of the decay based on the mass difference of the initial and final states, or Q-value. In the rest frame of the neutron, this energy from beta decay will be carried away predominantly by the electron and electron antineutrino since the proton is significantly more massive.

Experiments involving beta decay have been important in the development of nuclear physics. In 1957, Wu [8] measured that beta decay in nuclei of ^{60}Co violated parity symmetry based off of an experimental design developed in [9]. This experiment showed that only particles (antiparticles) with negative (positive) helicity, which is the projection of their spin along the direction of motion, interact via the weak force. Because the direction of a particle's spin is constant with respect to a boosted reference frame while the direction of the momentum of a massive particle is not constant in all boosted reference frames, neutrinos have zero mass in the simplest versions of the Standard Model [10], although this would be challenged by future measurements of neutrino oscillations.

1.3 Neutrino Oscillations

In the 1960's Davis conducted his now famous Homestake experiment, where he used incident solar neutrinos on ^{37}Cl nuclei to find that the total measured flux of neutrinos coming from the sun was only about a third of that expected from solar models [11]. It was Pontecorvo [12] who proposed the solution that neutrinos might oscillate from one flavor into another if the flavor states of neutrinos (governing their interactions with particles) were a non-trivial linear combination of the mass states (governing propagation in time). For a system of N different flavors and N different mass states we can use the unitary PMNS [13, 14] matrix U, described in greater detail at the end of this subsection, to show the relationship between the neutrino flavor states $\vec{\nu}_\alpha$ and the mass states $\vec{\nu}_j$ [15]

$$\vec{\nu}_\alpha = U\vec{\nu}_j,$$
$$\nu_\alpha = \sum_{j=1}^{N} U_{\alpha j}\nu_j.$$

Viewing this as a wavefunction $\psi_\alpha(t)$ as a function of time, we get

$$|\psi_\alpha(t)\rangle = \sum_{j=1}^{N} e^{-i\phi_j} U_{\alpha j} \, |\nu_j\rangle \tag{1.3}$$

with the standard relativistic phase factor $\phi_j = E_j t - \vec{p}_j \cdot \vec{x}$. Assuming a neutrino starting in the flavor state α, the probability of measuring a neutrino in flavor state β after some time t can be calculated via

$$P(\alpha \to \beta, t) = |\, \langle \nu_\beta | \psi_\alpha(t) \rangle \,|^2 \tag{1.4}$$

where

$$\langle \nu_\beta | = \sum_{j=1}^{N} \langle \nu_j | \, U_{j\beta}^{\dagger}. \tag{1.5}$$

In the simple case with just two neutrino flavors, the unitary matrix can be represented using a single parameter θ_{12}, the mixing angle between the two states. This representation yields

$$U = \begin{bmatrix} \cos(\theta_{12}) & \sin(\theta_{12}) \\ -\sin(\theta_{12}) & \cos(\theta_{12}) \end{bmatrix} \tag{1.6}$$

where a mixing angle of $\theta_{12} = 0$ yields the trivial result where the flavors aren't mixed, and $P(\alpha \to \beta, t) = \delta_{\alpha\beta}$. Expanding the oscillation probability for the case of an electron neutrino oscillating to a muon neutrino, for example, yields

$$\begin{aligned}
P(e \to \mu) =& | \langle \nu_\mu | \psi_e(t) \rangle |^2 \\
=& | \langle \nu_1 | \, (-\sin(\theta_{12})) e^{-i\phi_1} \cos(\theta_{12}) \, | \nu_1 \rangle + \langle \nu_1 | \, (-\sin(\theta_{12})) e^{-i\phi_2} \sin(\theta_{12}) \, | \nu_2 \rangle \\
& + \langle \nu_2 | \, (\cos(\theta_{12})) e^{-i\phi_1} \cos(\theta_{12}) \, | \nu_1 \rangle + \langle \nu_2 | \, (\cos(\theta_{12})) e^{-i\phi_2} \sin(\theta_{12}) \, | \nu_2 \rangle \, |^2 \\
=& | -\sin(\theta_{12}) e^{-i\phi_1} \cos(\theta_{12}) + \cos(\theta_{12}) e^{-i\phi_2} \sin(\theta_{12}) |^2 \\
=& 2\sin^2(\theta_{12}) \cos^2(\theta_{12}) - \sin^2(\theta_{12}) \cos^2(\theta_{12}) e^{-i(\phi_1 - \phi_2)} - \sin^2(\theta_{12}) \cos^2(\theta_{12}) e^{i(\phi_1 - \phi_2)} \\
=& \sin^2(\theta_{12}) \cos^2(\theta_{12}) (2 - e^{-i(\phi_1 - \phi_2)} - e^{i(\phi_1 - \phi_2)}) \\
=& \sin^2(\theta_{12}) \cos^2(\theta_{12}) (2 - 2\cos(\phi_1 - \phi_2)) \\
=& 4\sin^2(\theta_{12}) \cos^2(\theta_{12}) \sin^2\left(\frac{\phi_1 - \phi_2}{2}\right) \\
=& \sin^2(2\theta_{12}) \sin^2\left(\frac{\phi_1 - \phi_2}{2}\right)
\end{aligned}$$

Noting that $\phi_j = E_j t - p_j x$, and assuming relativistic neutrinos such that $p_j = \sqrt{E_j^2 - m_j^2}$ and $t = x = L$ where L is the baseline traveled by the neutrino from source to detection, this gives

$$\phi_1 - \phi_2 = \frac{m_1^2 L}{2E} - \frac{m_2^2 L}{2E} \tag{1.7}$$

Assuming the neutrinos are at the same energy. Putting this all together yields

$$P(e \to \mu) = \sin^2(2\theta_{12}) \sin^2 \left(\frac{\Delta m^2_{12} L}{4E} \right) \tag{1.8}$$

where Δm^2_{12} is the difference of the square of the masses. Assuming units for m_{12} in eV, L in m, and E in MeV, as well as Natural units, this simplifies to

$$P(e \to \mu) = \sin^2(2\theta_{12}) \sin^2 \left(1.27 \Delta m^2_{12} \frac{L}{E} \right) \tag{1.9}$$

This yields the standard formula for neutrino flavor oscillations.

While this was theoretically possible, there was no reason *a priori* that any of the masses or mixing angles between neutrino flavors should be non-zero. It was the work of the SNO [16, 17] and KamLAND [18] experiments that first measured solar neutrino oscillations by measuring rates of electron neutrino charged-current interactions alongside neutral-current interactions, which interact for all neutrino flavors. By showing that neutral-current interactions were consistent with solar models (SNO) and charged-current interactions were consistent with the Homestake results (SNO and KamLAND), the two experiments showed strong evidence for solar neutrino oscillations and non-zero neutrino oscillation parameters. Similarly, the Super-K experiment showed evidence for oscillations of atmospheric neutrinos [19].

Ultimately this led to the effort to characterize neutrino oscillations, which can be parameterized by three independent rotation matrices with three mixing angles and three mass squared splittings as formulated in Eq 1.9 (the CP violating phase and two phases associated with a Majorana mass term are ignored for this discussion) in the PMNS neutrino mixing matrix [13, 14]. These parameters have since been experimentally measured, and are presented in Table 1.1 [20].

Angle	Value	Error	Mass Split	Value [eV^2]	Error	Reference
$\sin^2\theta_{12}$	0.307	0.013	Δm^2_{21}	7.53e-5	1.8e-6	KamLAND + Solar
$\sin^2\theta_{23}$	0.545	0.021	Δm^2_{23}	2.453e-3	3.4e-5	Atm, Reactor, and Acc
$\sin^2\theta_{13}$	0.0218	0.0007	Δm^2_{13}	$\sim \Delta m^2_{23}$	3.4e-5	Reactor

Table 1.1: Measured mixing parameters for 3-flavor neutrino oscillations (assuming normal mass ordering) based on PDG data. θ_{12} and Δm^2_{21} come from data predominantly from KamLAND data with global fits from solar neutrino experiments [21]. θ_{23} and Δm^2_{23} comes from measurements across atmospheric, reactor, and long-baseline accelerator experiments [22]. θ_{13} measurements come from reactor experiments, particularly [23–26]

1.4 Neutrino Mass

In the Standard Model neutrinos have zero mass, but oscillation experiments directly measure the difference of the squares of $\bar{\nu}$ mass states. Oscillations therefore imply that at least two neutrino mass states are non-zero. Limits from experiments measuring deviations in the measured endpoint of beta decay spectra like KATRIN [27] and Project 8 [28] have set upper limits on the effective mass of $<$1.1 eV (90$\%$ CL) [29], with future results expected to probe sensitivities down to $\sim$40 meV [30].

This, coupled with the unnatural relative mass difference between neutrino masses and those of other fundamental particles (at least six orders of magnitude), has led to possible extensions to the Standard Model to account for neutrino masses. Although only left-handed neutrinos interact via the weak force, right-handed neutrinos could still exist. This could allow a Yukawa coupling, incorporating a Dirac mass term D after symmetry breaking of the Higgs field. On its own this could generate the neutrino mass, although it wouldn't explain the extremely small masses. Additionally, because of the neutral electric charge of neutrinos, there could be Majorana mass terms for either left or right handed neutrinos (m_L and M_R respectively). These mass term contributions to the Standard Model lagrangian, if they exist, would take the form

$$\mathcal{L}_\beta = \frac{1}{2} m_{L,\beta} \chi^\alpha \chi_\alpha + D_\beta \chi^\alpha \eta_\alpha + \frac{1}{2} M_{R,\beta} \eta^\alpha \eta_\alpha \tag{1.10}$$

where χ is a left-handed neutrino spinor, η is a right-handed neutrino spinor, and β iterates over the three neutrino flavors. Rewriting this in matrix form yields

$$\mathcal{L} = \begin{bmatrix} \chi & \eta \end{bmatrix} \begin{bmatrix} m_L & D^T \\ D & M_R \end{bmatrix} \begin{bmatrix} \chi \\ \eta \end{bmatrix} \tag{1.11}$$

Here we see the mass matrix, and can derive some properties of the eigenvalues (the relevant terms regarding measured neutrino mass) if we include some simple assumptions. First assume that the Dirac masses D are of the same scale as the other fundamental particles, around 10^{11} eV. Let us also assume that the left-handed majorana mass is much smaller, as constrained by some neutrinoless double beta decay experiments [31, 32]. Finally, the right-handed majorana mass term can be arbitrarily large. Some [33] even suggest it could be on the GUT scale (10^{25} eV). Ultimately this leads to the general assumption that $m_L \ll D \ll M_R$. under these assumptions, to leading order the two eigenvalues are M_R, which is beyond the current reach of experimental measurements, and $m_L - \frac{D^2}{M_R}$.

This latter eigenvalue, whether m_L is small or zero, leads to a measured mass on the sub eV scale in line with current limits, while both connecting to the weak scale of other fundamental particles and GUT energy scales. In the case where m_L is identically zero, this is called the type-I seesaw mechanism (m_L dominating over $\frac{D^2}{M_R}$ is referred to as the type-II seesaw mechanism) [33] because the heaviness of the one eigenvalue lightens the other eigenvalue (if one goes down, the other goes up). This promising mechanism for neutrino mass generation is yet to be verified experimentally, does not have any bearing on the results of this work, and will not be discussed further in this manuscript.

1.5 Reactor Antineutrino Anomalies

Besides solar and atmospheric sources, nuclear reactors have historically provided a rich test bed for neutrino physics as well. Neutrinos produced by nuclear reactors were used in the first experimental measurement of neutrinos [5], and they were used by KamLAND to measure the mixing parameter θ_{12} [18]. The cores of nuclear reactors contain fissile isotopes. The nuclei split into neutron rich isotopes that then beta decay to new isotopes which continue the chain of beta decays themselves. In equilibrium, reactors can contain thousands of different nuclear isotopes and produce large fluxes of antineutrinos. Pressurized water reactors are commonly used for electrical power generation, and contain low enriched uranium cores (LEU). These contain a mixture of different fissile isotopes, and as a reactor core burns over time those relative fission fractions will evolve as one isotope burns up and feeds another isotope, such as the burning of ^{235}U producing ^{239}Pu. This usually leads to a time-dependence of the neutrino production in a reactor if there are no adjustments made from an operations standpoint to mitigate this effect. Many experiments have measured both the flux and spectrum of antineutrinos produced in nuclear reactors, and those measurements have led to two major "anomalies" regarding the difference between measurements and model predictions.

1.5.1 Reactor Flux Anomaly

In addition to the solar neutrino flux anomaly associated with the measurement from the Homestake experiment, recent experiments have demonstrated discrepancies between measurements and predictions in other sectors of neutrino physics. A global neutrino flux analysis [34] found a global rate deficit of short baseline ($<$100m) reactor antineutrinos. Additionally, an excess of $\bar{\nu}_e$ events in the LSND and MiniBoone experiments [35, 36] and a deficiency in ν_e events in the GALLEX and SAGE experiments [37] have been

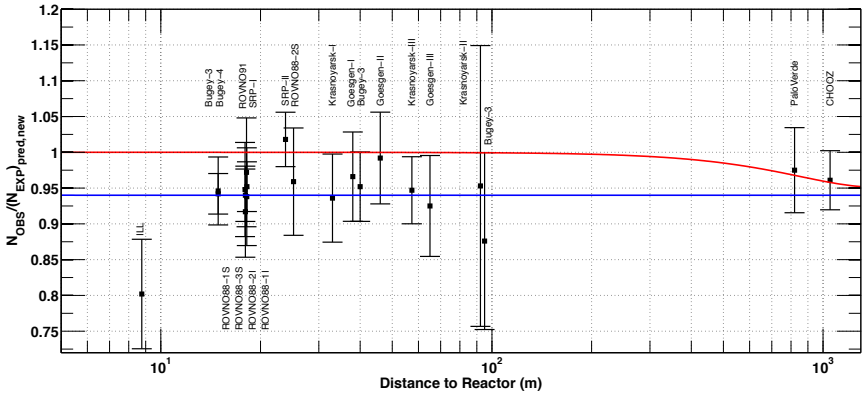

Figure 1.2: Illustration of Reactor Antineutrino Anomaly (RAA) from global analysis [34] looking at the ratio of measured reactor antineutrino flux to expected flux with no oscillations for several experiments. The red line shows a best fit with the introduction of 3 standard neutrino oscillations while the blue line shows predictions with a fourth neutrino state with $\Delta^2_{new} \sim 1eV^2$ for illustration purposes.

measured. These measurements across a wide range of energies, baselines, and complementary experimental strategies incorporating reactor antineutrinos, accelerator neutrinos, and neutrinos from radioactive sources all point towards a possible new neutrino with a mass splitting squared $\sim$1eV^2, as suggested in [34]. This additional neutrino state could allow for oscillations between the ν_μ and ν_e flavor states on the appropriate $\frac{L}{E}$ scale to resolve these various anomalies much like neutrino oscillations of solar neutrinos resolved the Homestake anomaly, as shown for the reactor neutrino case in Fig 1.2. From [34], a best fit point of sin$^2(2\theta_{14})$ = 0.23 and Δm^2 = 1.88eV^2 is found, commonly referred to as the Reactor Antineutrino Anomaly (RAA) Best Fit Point. An active neutrino with this mass splitting, however, is ruled out by the measurement of the width of the Z-boson decay [38], which strongly limits the number of weakly interacting neutrino flavors with masses less than half the mass of the Z-boson to 3. Thus a new neutrino flavor with mass splitting squared of $\sim$1eV^2 could only be non-weakly interacting, or a sterile neutrino.

1.5.2 Reactor Spectrum Anomaly

In addition to the total flux anomaly, deviations in the spectrum shape of reactor antineutrinos were originally measured by the Daya Bay [39, 40], Double Chooz [41], and RENO [42] experiments. All of these measurements found general disagreement between the measured spectrum of low enriched uranium (LEU) reactor antineutrino spectra from model predictions, with the most notable deviation being a relative excess in the 4-6 MeV prompt energy range (shown in Fig 1.3).

This is problem is further complicated by the composition of the LEU reactor cores measured by these experiments. Having an admixture of four different fissile isotopes (^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu), it is difficult to deconvolve the contributions towards this excess from each isotope. It may be the case that a single isotope is responsible for this discrepancy, they could all be contributing equally, or there could be a non-trivial linear combination of contributions to this feature.

1.5.3 Model Predictions

Model predictions for reactor antineutrino flux are generated in two main ways. The *ab initio* method (such as what's described in [43]) uses nuclear databases to directly sum up the expected contributions from thousands of individual decays associated with the beta decays of fission fragments of an isotope and their decay products in a nuclear reactor. The total energy spectrum S as a function of neutrino energy $E_{\bar{\nu}}$ is calculated by

$$S(E_{\bar{\nu}}) = \sum_{i=0}^{n} R_i \sum_{j=0}^{m} f_{ij} S_{ij}(E_{\bar{\nu}}) \tag{1.12}$$

where R_i is the equilibrium decay rate of isotope i, f is the branching fraction of the ith isotope decay to the j energy level, and S_{ij} is the individual energy spectrum for each isotope to each energy level. Of the thousands of individual spectra used to calculate the

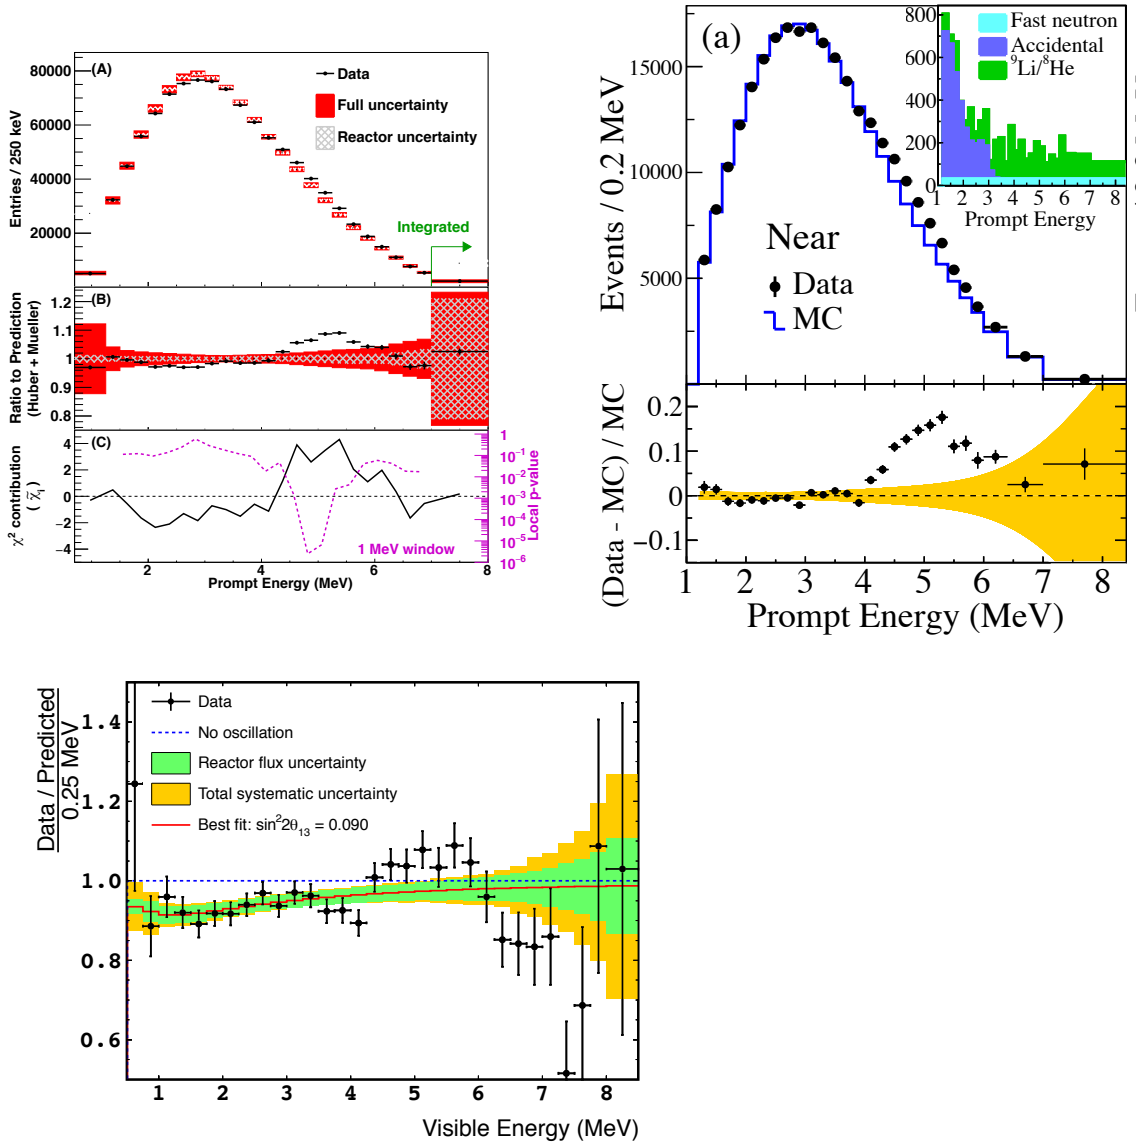

Figure 1.3: Reactor antineutrino spectrum measurements from (left) Daya Bay [40], (right) Reno[42], and (bottom) Double Chooz [41]. The ratios to predictions all show a significant excess in the 4-6 MeV range.

total antineutrino spectrum for a given parent fissile isotope, roughly 5% of the neutrino yield decays are missing. This makes it extremely difficult to quantify the uncertainties properly for this method.

A second, independent method used to predict the total antineutrino spectrum for a given isotope is the beta-conversion method done by Ref. [44]. This method uses direct measurements of the total beta decay spectrum from irradiated foils of fissile isotopes at ILL [45–47]. Then by fitting virtual beta branches to the data, corresponding antineutrino spectra can be converted from each beta branch. This implements assumptions that the transitions are all allowed transitions, and limits the shape effects that come from forbidden transitions. This beta-conversion model, most recently implemented by Huber [44] and sometimes called the Huber-Mueller Model (using Mueller's version of the ^{238}U spectrum [48]), has been the commonly accepted best prediction for the antineutrino spectrum for the fission of ^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu. [49] shows the converted reactor antineutrino spectrum from two implementations of this method with different assumptions regarding subdominant corrections in Fig 1.4.

Whether the antineutrino spectrum is produced using an *ab initio* method or the beta-conversion method, it must be convolved with the IBD cross section [50] in order to predict the measured spectrum as presented in Figure 1.3. In [49] the cross section to zeroth order in $1/M$, where M is the mean mass of the proton and neutron in the IBD interaction, is given as

$$\sigma_{tot}^{(0)} = \frac{2\pi^2/m_e^2}{f^R \tau_n} E_e^{(0)} p_e^{(0)} \tag{1.13}$$

where m_e is the mass of the electron, f^R is the neutron decay phase space factor, τ_n is the lifetime of the neutron, and $E_e^{(0)}$ and $p_e^{(0)}$ are the zeroth order estimates of the energy and momentum of the outgoing positron. Since reactor antineutrinos are at relatively low energy, [49] further includes the first order energy term, and also introducing radiative

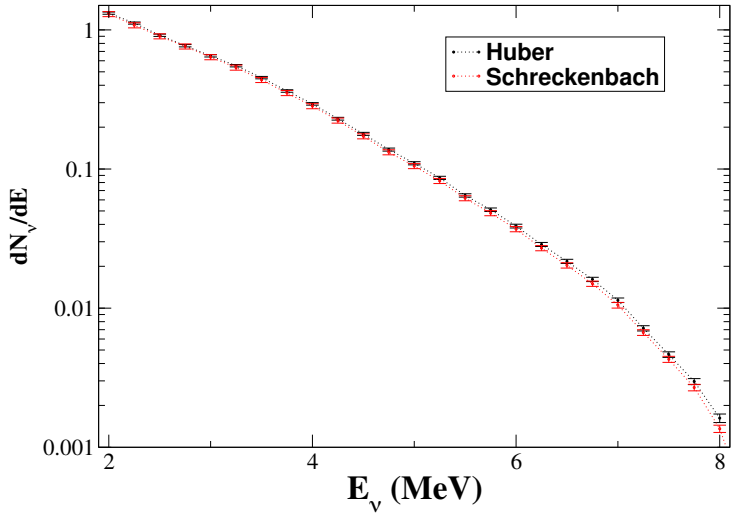

Figure 1.4: Two calculations of the emitted antineutrino spectrum for the fission of ^{235}U using the beta-conversion method. The difference arises from differences in the assumptions made about the subdominant corrections to beta-decay [49].

corrections of order $\frac{\alpha}{\pi}$ yields a simplified form of

$$d\sigma(E_\nu E_e) \rightarrow d\sigma(E_\nu E_e)\left[1+\frac{\alpha}{\pi}\left(6.00+\frac{3}{2}\log\frac{M_p}{2E_e}+1.2\left(\frac{m_e}{E_e}\right)^{1.5}\right)\right] \tag{1.14}$$

A numerically calculated plot of both the zeroth order and first order total IBD cross section is shown in Fig 1.5 from [51]. This plot also includes the expected relative scattering angle of the positron for reference. Also note that this plot showcases the absolute threshold E_ν^{thresh} for IBD to occur dictated by the final state mass and recoil constraints

$$E_\nu^{\text{thresh}} = \frac{(M_n+m_e)^2-M_p^2}{2M_p} = 1.806 \text{ MeV} \tag{1.15}$$

The final prediction of the measured reactor antineutrino spectrum via IBD interactions (as is used in Figure 1.3) is generated by convolving the reactor spectrum with the IBD cross section.

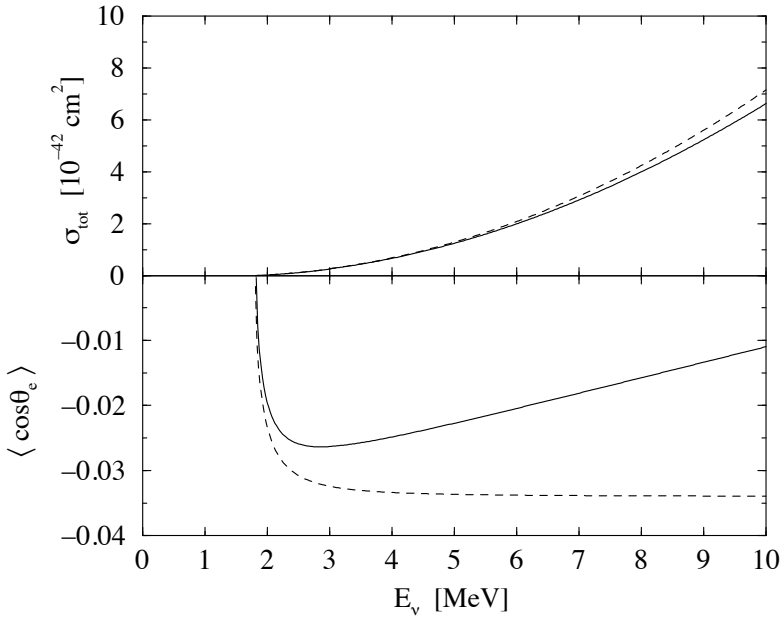

Figure 1.5: (top) Total IBD cross section and (bottom) expected positron scattering angle as a function of neutrino energy. The dashed (solid) lines give the zeroth (first) order calculation in $\frac{1}{M}$ [51].

1.6 Motivation for New Precision Measurements

For the anomalies outlined in Section 1.5, there are a number of possibilities to explain the discrepancy between the measurements and model predictions, both for the integrated rate and spectral shape. Concerning the rate deficiency, sterile neutrinos could explain the deficit as could incorrect predictions of the total flux of at least one contributing isotope. The spectral shape discrepancy could be the result of some correlated systematic effect in the prediction of the reactor spectra, contributing to deviations in all of the input spectra, or it could be an issue with the prediction of a singular fission isotope. In order to resolve these discrepancies, new precision measurements are needed of reactor antineutrino spectra. It is with this in mind that the PROSPECT experiment was designed to measureof the antineutrino spectrum of a highly enriched uranium reactor at multiple baselines. This enables the search for sterile neutrino oscillation signatures while also making a precision

measurement of the antineutrino spectrum of ^{235}U. The PROSPECT detector design, as well as the analysis of its 2018 data run, are detailed in the following chapters of this thesis. Further, a joint analysis of the PROSPECT and Daya Bay spectrum measurements is described and improved results are presented.

Chapter 2

PROSPECT Detector Design

2.1 Overview of PROSPECT and Design Goals

It is with the goals described in Section 1.6 in mind that the Precision Reactor Oscillation and Spectrum (PROSPECT) [52] experiment was designed and deployed at the High Flux Isotope Reactor [53] at Oak Ridge National Laboratory [53] as shown in Figure 2.1. PROSPECT consists of 4 tons of a ^{6}Li-doped liquid scintillator (LiLS) optically segmented into 154 discrete cells that detect reactor antineutrinos via IBD events [54]. Photomultiplier tubes (PMT) are located on the two ends of each segment, and the entire optical lattice and PMT structure are supported by an acrylic structure in an inner acrylic vessel inside of a larger aluminum tank and shielding package. The entire detector is characterized in part using a radioactive source calibration system and also partially characterized using cosmogenic sources. Each of these components, as well as the reactor source, is described in more detail later in this section. A drawing showing many of these components is given in Figure 2.2 [55].

In the PROSPECT detector, a prompt flash from the positron annihilating with an electron closely followed by the signal of a neutron capturing on ^{6}Li is used to identify

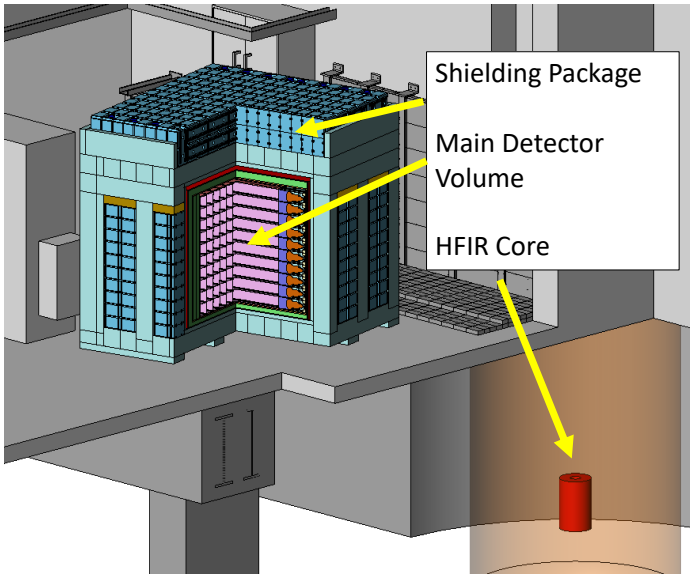

Figure 2.1: General layout of the PROSPECT detector at HFIR. The main detector volume, as well as the shielding package are situates above grade from the reactor core at an average baseline of approximately 8 m.

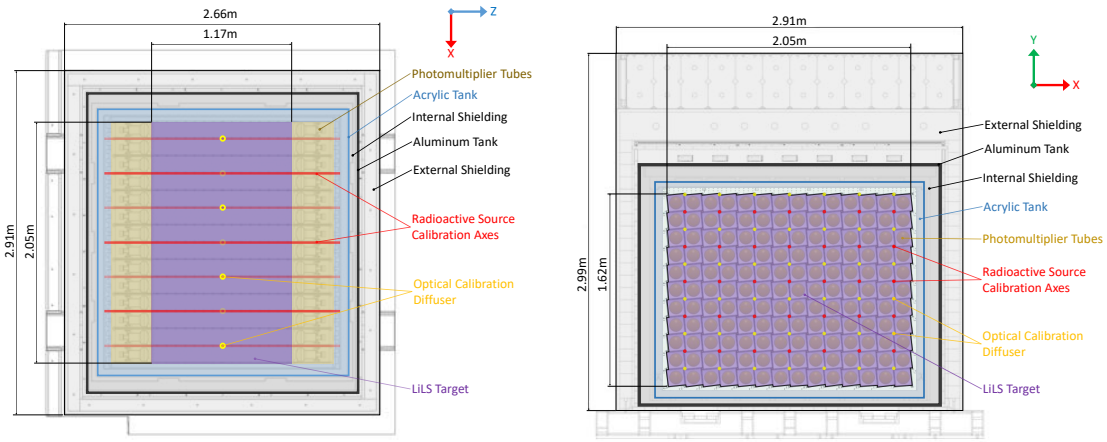

Figure 2.2: A top (left) and side (right) view drawing of the PROSPECT detector main features [55].

IBD-like events.

$$n + {}^{6}\mathrm{Li} \rightarrow \alpha + t \tag{2.1}$$

Equation 2.1 shows how the neutron produced in the LiLS captures on ^{6}Li to produce an alpha and a triton particle, which deposit energy in a localized spot in the detector. The detector was deployed to cover a range of baselines from 7-9m. This short baseline coverage allows PROSPECT to probe higher Δm^2 regions in phase space than many other experiments; additionally the close proximity to the reactor yields a relatively high flux of neutrinos through the detector [56].

In order to make a meaningful measurement towards searching for sterile neutrinos and determining the antineutrino spectrum of ^{235}U, PROSPECT was designed with several goals [56]:

- An energy resolution better than $10\% \ / \sqrt{E(\mathrm{MeV})}$ (RMS)
- A position resolution better than 20cm
- A signal-to-background ration greater than 1:1
- A detector mass of a few tons with a baseline coverage of $\sim$ 3m

2.2 Detector Design

These goals were all either met or exceeded in the experiment, with technical descriptions following for the HFIR source as well as the production and installation of the LiLS scintillator, optical separators, PMT optical modules, acrylic containment, aluminum containment, shielding, and calibration systems.

2.2.1 HFIR

PROSPECT is deployed at the High Flux Isotope Reactor (HFIR) at Oak Ridge National Laboratory (ORNL) [53]. The HFIR reactor [57] is a light water, research reactor that has a core with two, concentric rings of ^{235}U (enriched at 93% by mass), and is one of the few highly enriched uranium (HEU) reactors in the US. The highly enriched fuel, compact core, and relatively short reactor cycles of HFIR makes it an ideal source of reactor antineutrinos for the PROSPECT experiment [55]. Modeling and simulation of the core across a full reactor cycle [58] show that the fission rate change is negligible across the cycle and that the ^{235}U fission fraction remains over 99.5% throughout the cycle. This ensures that the underlying reactor antineutrino spectrum measured by PROSPECT does not change with time across each reactor cycle. Additionally the core of HFIR is relatively compact. The cylindrical core measures 50.8cm tall with a diameter of 43.5cm [58], which helps to reduce the uncertainty of the antineutrino baseline L used in the oscillation analysis (discussed in greater detail in Section 4). Figure 2.3 shows a side view drawing of the reactor. The figure also shows top and side view photographs of a dummy HFIR fuel element, a representation of the different components modeled in a MCNP model [59] distinguished by color, and a projection of the fission power density during operation. The final major feature of HFIR relevant to PROSPECT is its reactor cycle schedule. HFIR nominally operates at 85 megawatts thermal in cycles that last 23-26 days at approximately 7 cycles per year. This corresponds to roughly 46% reactor up time across the year. Besides the short reactor cycles contributing to the stability of the ^{235}U fission fraction, having roughly equal time reactor on and reactor off allows PROSPECT to directly measure environmental backgrounds and subtract them off [58].

The PROSPECT detector was deployed in a passageway immediately adjacent to the reactor pool wall above grade of the core. The detector covered a baseline range of 6.7 to 9.2 meters from the core and had $<$1 meter water equivalent overburden. Original design

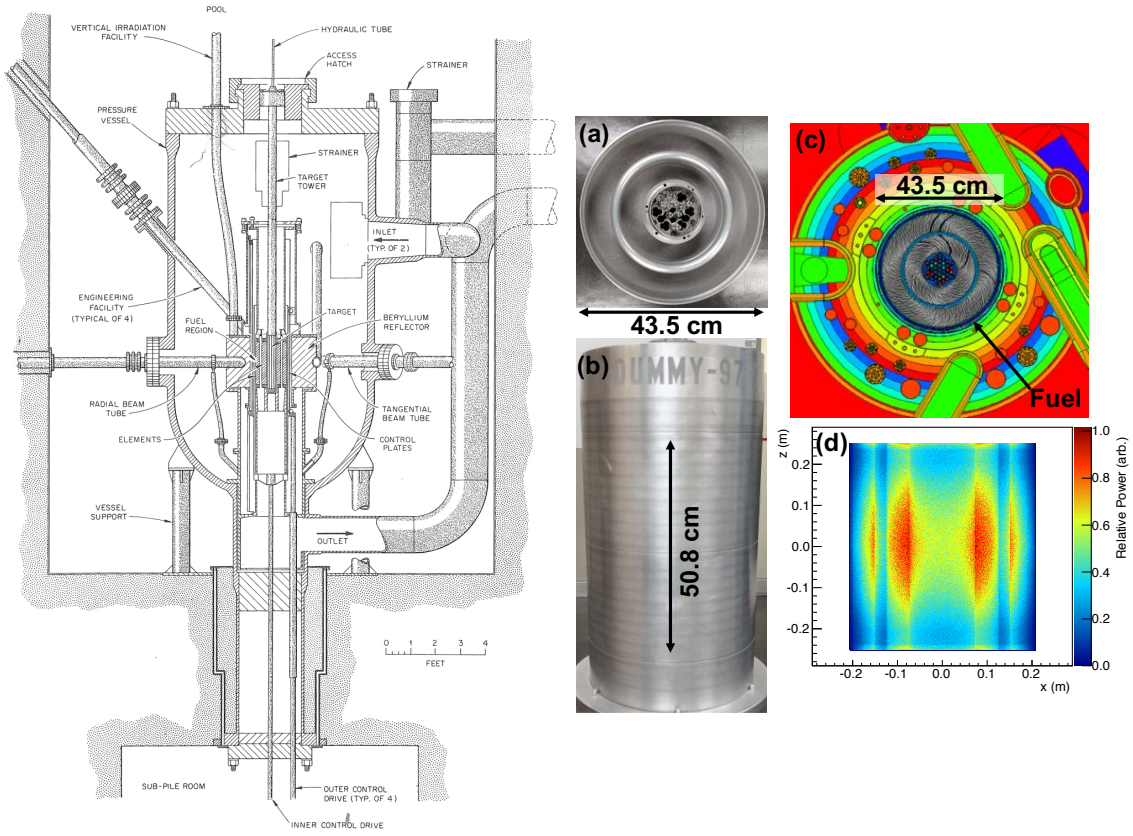

Figure 2.3: (left) Side view of HFIR reactor (right) (a,b) Dummy HFIR fuel element from the top and side. (c) Components of the Monte Carlo N-Particle (MCNP) model of the HFIR core. (d) 2D projection of the fission power density of HFIR [58].

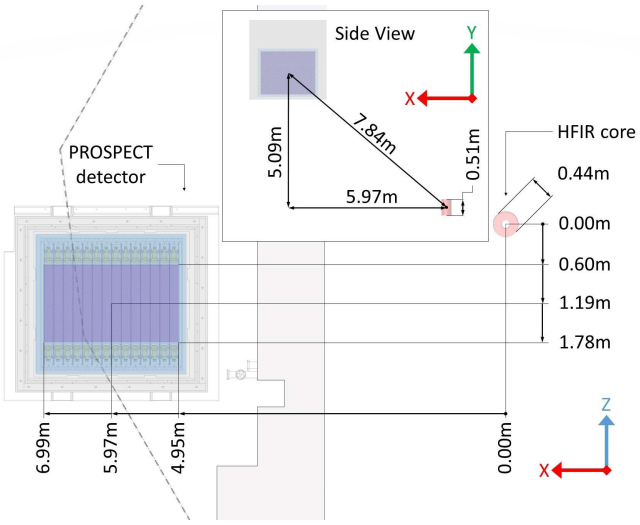

Figure 2.4: Top and side view of the PROSPECT detector layout with respect to the HFIR core, where the center-to-center distance between the two is approximately 7.84 m. Coordinate axes are given, where the xz-plane is parallel to the floor and $+y$ is the upward direction [58].

plans included for movement of the PROSPECT detector throughout its data taking, but due to logistical constraints this system was only used for moving the detector into place and not for moving throughout operation. The layout of the detector with respect to the HFIR core is given in Figure 2.4 [58]. Further details of the shielding package and passive background mitigation is covered in Section 2.2.6.

2.2.2 ^{6}Li-loaded Liquid Scintillator

The scintillator used for the PROSPECT experiment consists of a base of a commercial, di-isopropylnaphthalene-based scintillator (EJ-309) [60]. Besides the addition of 2,5-diphenyloxazole (PPO) and 1,4-bis(2-methylstryryl) benzene (bis-MSB) acting as a wavelength shifters and an ether-based glycol nonionic surfactant, a 9.98 mol/L aqueous solution of LiCl with 95% enriched ^{6}Li by atom [61]. This ^{6}Li acts as the main thermal neutron capture agent in the scintillator in addition to hydrogen. Additionally, a small but precise amount of ^{227}Ac is spiked into the scintillator for calibration purposes [62].

The scintillator was prepared at Brookhaven National Lab (BNL) using samples of the

enriched LiCl provided by the National Institute of Standards and Technology (NIST). Altogether, twenty-eight 55 gallon drums were shipped from BNL to Oak Ridge, where all drums of LiLS were pumped into a 24,000 L capacity ISO tank. Nitrogen gas lines were installed to allow for mixing of the LiLS to ensure uniformity before pumping approximately 4521 L of LiLS into the detector.

Samples of the LiLS were taken both before shipping to ORNL as well as periodically throughout the detector filling process for quality assurance. A histogram of the measured light yield of the LiLS production batches is included in Figure 2.5 [61]. The light yield was measured using the ratio of the Compton edge of a measured ^{137}Cs spectrum to the same feature in a linear alkylbenzene (LAB) reference scintillator. The acceptance criterion was set to 95% of the light yield of the LiLS used in the PROSPECT-50 prototype [63].

Besides the high light yield, an important feature of the LiLS is the ability to reject backgrounds through the use of particle identification via pulse shape discrimination (PSD). The energy deposition of the smaller, electron-like events differs from those of the heavier, nuclear recoil-like events by producing light pulses with different lifetimes. This leads to pulses of light with different shapes for different classifications of events. A PSD metric is defined as the integrated charge in the tail of the waveform q_{tail} divided by the total integrated charge of the waveform q_{tot}

$$PSD = \frac{q_{tail}}{q_{tot}}. \tag{2.2}$$

In this equation, q_{tail} is integrated from 44 to 200 ns after the leading edge of the pulse while q_{tot} is integrated from 12 ns before to 200 ns after the leading edge [55]. From fitting two gaussians to a distribution of PSD metrics from two types of events (electron-like and nuclear-like) like that in Figure 2.6 [61], a figure of merit (FOM) is defined to quantify

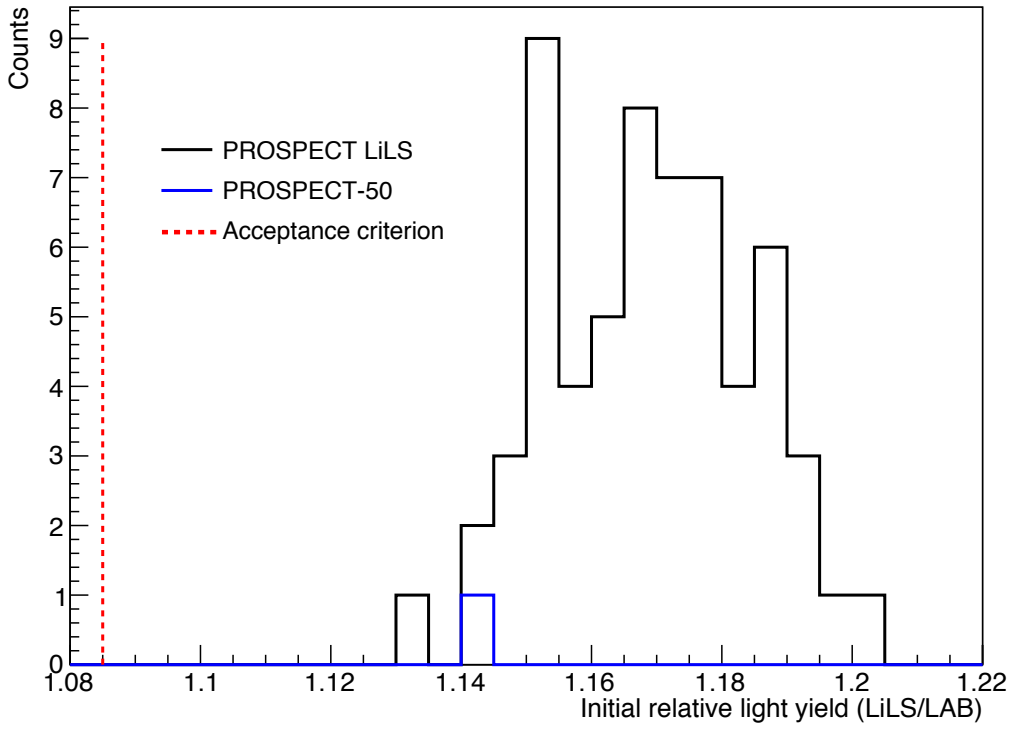

Figure 2.5: Histogram of initial measured sample light yields of LiLS relative to LAB reference sample. The acceptance criterion (red dashed line) was set to 95% of the measured light yield in the PROSPECT-50 prototype [61].

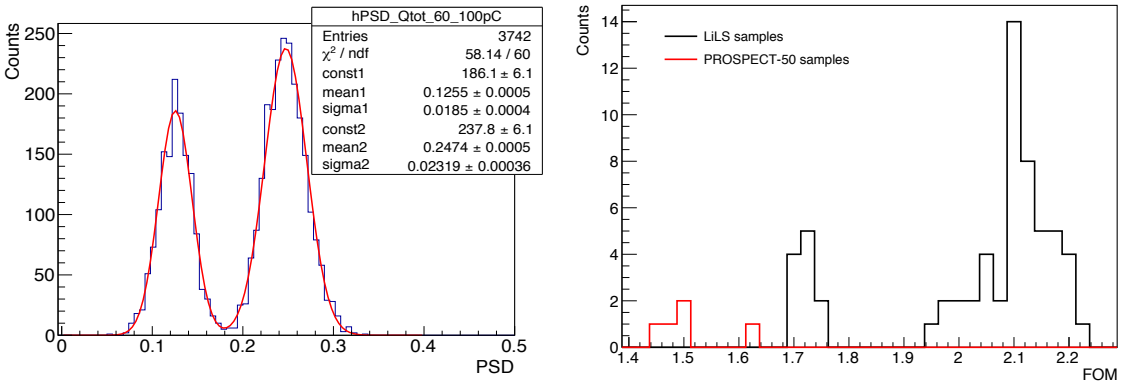

Figure 2.6: (left) Example fitted PSD distribution for total detected events from a ^{241}Am-^{9}Be neutron around 540 keV visible energy showing a high figure of merit to distinguish event types. (right) Distribution of FOM (n, ^{6}Li) for the LiLS batches, along the measurements from the PROSPECT-50 prototype. The two modes in the LiLS samples come from a procedural improvement introduced partway through production [61].

the level of event identification:

$$FOM = \frac{\mu_n - \mu_e}{\sqrt{FWHM_n^2 + FWHM_e^2}}. \tag{2.3}$$

Here μ_n (μ_e) is the fitted mean of the nuclear-like (electron-like) distribution and $FWHM_n$ ($FWHM_e$) is the full width at half the maximum of the nuclear-like (electron-like) distribution. An example fitted distribution from one of the LiLS samples is given in the left panel of Figure 2.6 [61] while a distribution of the FOM values of all LiLS batches is included in the right panel. FOM values from the PROSPECT-50 prototype are included for reference in red, and all accepted batches have a FOM value higher than what was measured in the prototype measurement. The two modes in the LiLS distribution arise from a procedural change that was introduced partway through LiLS production, leading to an increase in performance. All batches of LiLS were mixed before filling the detector to ensure uniform performance throughout the detector.

2.2.3 Optical Lattice

Design Goals

The PROSPECT optical grid partitions the full $2.045 \times 1.176 \times 1.607 \text{m}^3$ target volume into an 11×14 array. It primarily consists of highly reflective segment separators fixed in place using PLA rods, which interfaced directly with the PMT housings (a basic layout is shown in Figure 2.7 [64]). Each individual segment is $1.176 \times 0.145 \times 0.145 \text{m}^3$ in volume and is rotated 5.5° from the horizontal both so that the PLA rods can support them and to interface with the radioactive source deployment system.

The lattice was designed with several constraints in mind. The components must have high reflectivity for optical photons in the 400 nm to 550 nm range for efficient transport of scintillation light. Optical simulations and results from prototype detectors demonstrated that specular reflectors were more efficient than diffuse reflectors for this purpose [61]. Additionally the components must be opaque to eliminate optical crosstalk between segments of the detector. Components must also have as minimal mass and volume to reduce the amount of non-scintillating volume in the detector, as this effects the total energy that can be detected of the IBD positron.

The optical grid itself must be mechanically stable enough to handle the vibrations associated with shipping the detector between its locations of assembly and deployment, and its components must have a high degree of uniformity to ensure consistent segment performance. The optical grid itself must also properly interface with other components of the detector, namely the radioactive source deployment system as well as the LiLS. That is any component in direct contact with the LiLS must be chemically compatible with the liquid.

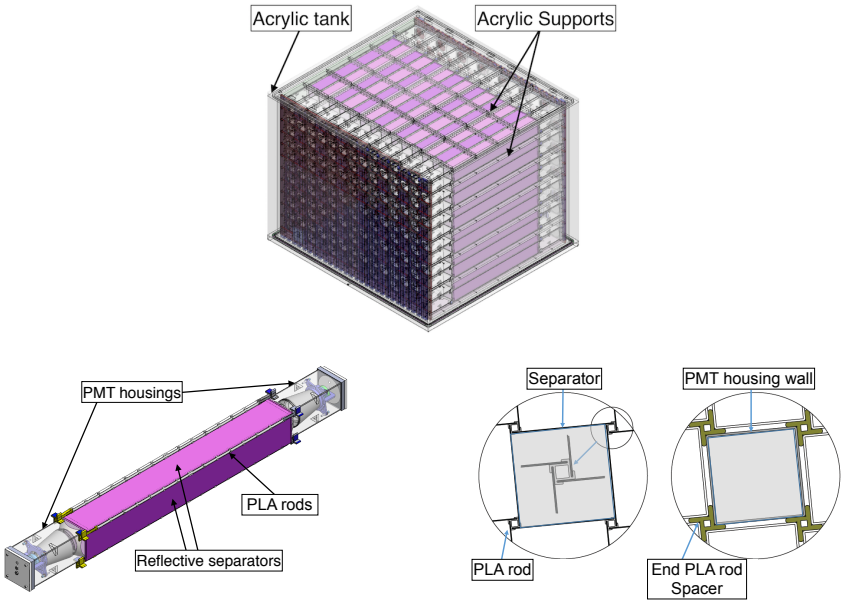

Figure 2.7: (top) Active detector volume enclosed in an acrylic support structure consisting of a large tank and supports. (left) A single optical segment consisting of four reflective separators supported by four PLA rods as well as two PMT optical modules on the ends. (right) cross section of the segment both in the longitudinal middle of the segment as well as at the interface with the PMT housing [64].

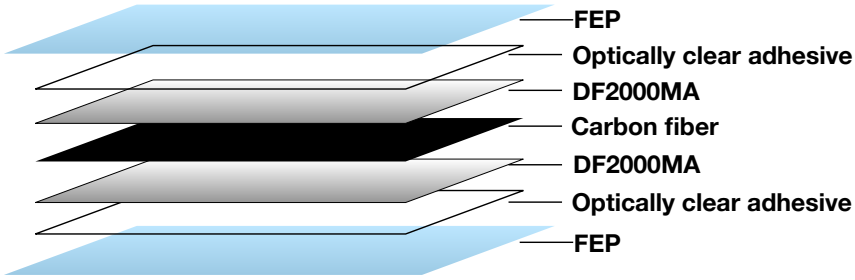

Figure 2.8: Layers of the optical separators. The inner layer is carbon fiber used for mechanical rigidity, followed outward by the DF2000MA reflective material, a transparent adhesive, and a sealed FEP outer layer to ensure chemical compatibility with the LiLS [64].

Reflective Separators

The separators themselves consist of a sandwich of materials. Working from the center outwards, the inner-most layer is a carbon fiber coated in an epoxy resin backing which provides the mechanical support for the separator. Overtop of this is the reflective material DF2000MA which is made from multiple layers of an organic polymer with multiple total internal reflections leading to an increased overall reflectivity [65]. Following that is a layer of a dual-sided adhesive and then an outer layer of transparent FEP laminated over that. This FEP extends beyond the edges of the carbon fiber backing and reflective material to ensure no LiLS comes in contact with any of the inner materials. A schematic of this separator sandwich is shown in Figure 2.8 [64]. The total thickness of this sandwich is 1.18 ± 0.05 mm, which maintains the low thickness requirement to minimize energy loss from the positrons.

PLA Rods

The support rods for the optical lattice were fabricated from white-dyed PLA, or polylactic acid, 3D printed via Fused Deposition Modeling. These rods were designed to be low-mass and act as the interface between the reflective separators, the PMT modules, and the

acrylic supports of the detector. The rods are rigid, longitudinal tubes with rectangular cross-sections which were utilized for deployment of calibration sources. There are also tabs on all of the rods used to capture the reflective separators.

Each rod consists of nine individual PLA portions strung along either a PTFE tube (for radioactive source deployment rods), or a square acrylic rod (all others). Some of the square acrylic rods also featured a groove cut along its length for the optical calibration system. Both the source and optical calibration systems are described in greater detail in Section 2.2.7. Each rod consisted of a sequence of an end piece, three standard pieces, a central piece that could have a cutout for the optical calibration system, three more standard pieces, and another end piece. Each end piece style is dependent on the rod's intended location in the detector, with spacers on either four sides (inner detector rods), three sides (edge detector rods), or two sides (corner detector rods). Each end piece also has a left-handed and right-handed counterpart. Figure 2.9 [64] shows each of the different types as well as their configuration along a central segment of the detector.

2.2.4 Photomultiplier Optical Modules

Each segment of the detector has a photomultiplier (PMT) optical module on each of its two ends used to detect the scintillation light. Two types of PMTs were used in the detector. The inner volume was instrumented with Hamamatsu R6592 SEL PMTs, while the outermost side and top PMTs were ADIT Electron Tubes 9372KB PMTs (ETL). See Figure 2.10 for the layout [54]. Two different types of PMTs were used to minimize costs for the detector and maintain parallel production pipelines for deployment on a timely schedule. The Hamamatsu PMTs performed with better specifications and were reserved for the inner, fiducial volume of the detector with most non-fiducial segments instrumented by ETL tubes.

These PMT modules are designed with a UV-transparent acrylic window facing the

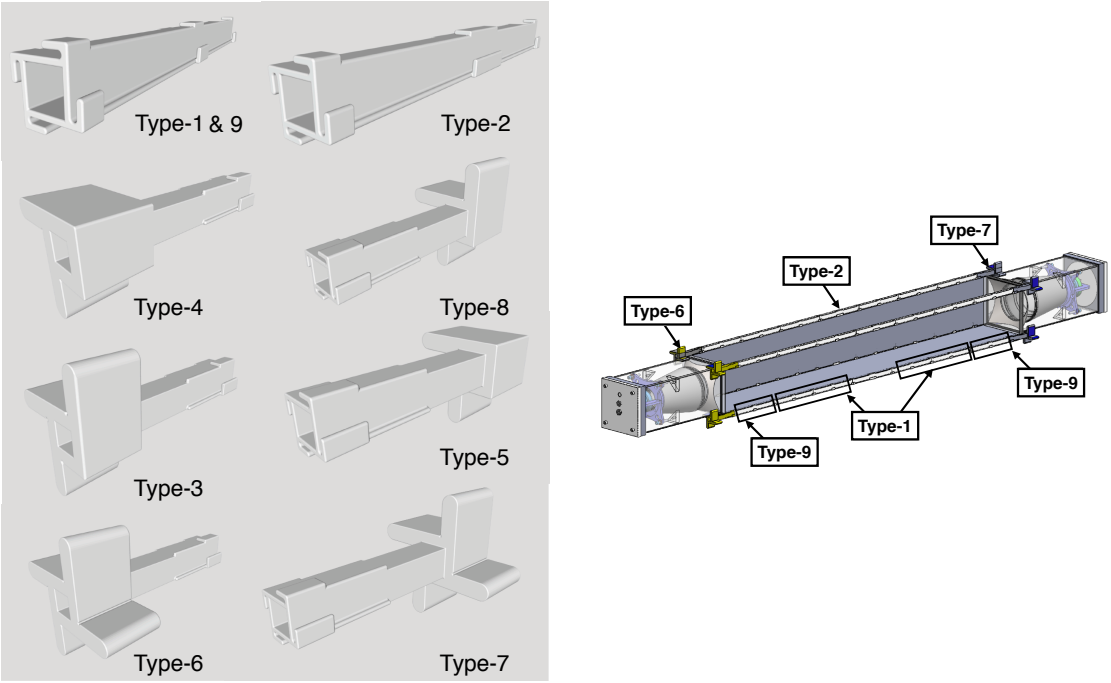

Figure 2.9: (left) Each of the different types of segments of the support rods for the detector. Types-1 & 9 are the standard pieces on each rod, Type-2 is the center piece, and Types 3-8 are the various configurations of end pieces. (right) A schematic of four support rods along the edges of a central optical segment of the detector [64].

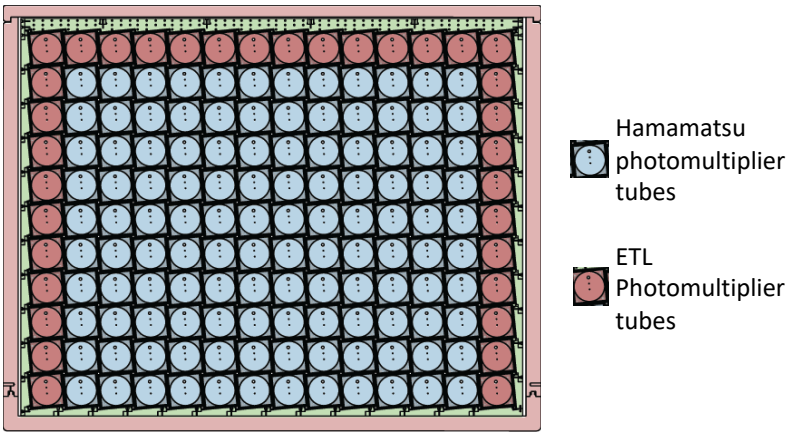

Figure 2.10: Cross-section of the PROSPECT detector. Squares in blue represent segments instrumented with Hamamatsu PMTs while the side and top segments had ET PMTs. Two PMT models were used to mitigate cost and timeline constraints while maintaining overall physics sensitivity, with non-fiducial segments mostly instrumented by ETL tubes [54].

active volume of the detector with feedthroughs for a signal and high voltage line out of its back. The back (sides) were constructed out of black (white)-dyed acrylic to reduce crosstalk between modules. Inside the module, in addition to the PMT, are several features. Laser-cut, acrylic PMT supports were glued in place and used stainless steel screws and springs to keep positive pressure of the PMT on a front facing reflector cone. Besides holding the PMT in place, this cone also functioned as a light guide to direct scintillation light onto the photocathode of the PMT. Additionally, a FINEMET [66] magnetic layer was placed around the main body of the PMT to shield effects from local magnetic fields. An acrylic plug was also attached in the back with two cable feedthroughs and a port. The cables feeding through the back plugs were soldered to printed circuit boards with tapered voltage dividers to maintain linear performance over a wide dynamic range of operation. These dividers were attached to the PMT pins themselves via a plastic socket. Through the port in the back plug, the entire PMT housing was filled with mineral oil to optically couple the PMT to the front face of the module, reduce the buoyant force on the module in the LiLS volume, and reduce the pressure differential experienced by the module. 150cc of air was included in small pouches that acted as an expansion volume for thermal fluctuations. These components, besides the air pouches, are shown in Figure 2.11 [54].

The containers are liquid tight so that the internal components never come in contact with the scintillating liquid. The only materials from the module that come in contact with the LiLS are the transparent, white-dyed, and black-dyed acrilic, PEEK feedthroughs for the cables, and Krytox grease and Viton O-rings used for the seals.

During production, all acrylic components were cleaned with a 1% solution by mass of Alconox and distilled water, while other components were cleaned with ethanol. All components were rinsed with distilled water until their runoff had no measurable conductivity (0-0.1 μS). Housings were qualified for use based off of measurements of the physical dimensions produced as well as liquid tightness of the boxes. Physical dimensions for PMT housings were also maintained in an SQLite database throughout the production process.

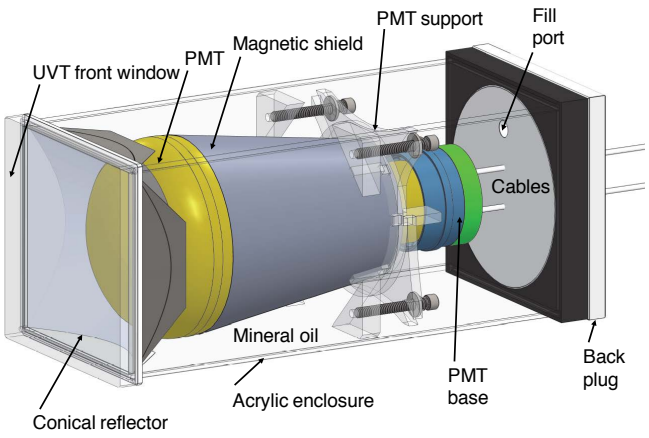

Figure 2.11: Internal components of a PMT optical module [54].

All PMTs also passed several electronic tests throughout the production stage. Two dark boxes were setup with a capacity for 16 PMTs each, which had cable feedthroughs and an LED pulser which could illuminate all PMTs at once. Figure 2.12 shows an open example setup of 16 PMTs.

A nominal voltage of 1600 (1150) was applied to the Hamamatsu (ET) PMTs for the initial test. An electronic trigger then was used to operate both the low amplitude LED pulser and a data acquisition setup for the PMTs. Integrating over the waveforms seen by each PMT, a distribution was made for the signals in integrated analog-to-digital (ADC) units. That distribution was then fit with a function consisting of the sum of an independent gaussian pedestal and 5 additional, correlated gaussians representing the single, double, third, fourth, and fifth photoelectron peaks. The single photoelectron mean was saved to an SQLite database for future analysis. Then, the PMTs were set to 1800V to burn in for 48 hours. Single photoelectron measurements were made again after this burn in to check for consistency of each PMT. An example comparison pre- and post- burn in is given for Hamamatsu PMT 68 in Figure 2.13.

Finally, the SPE measurement was done over a range voltages in 50V increments to

Figure 2.12: One of two dark boxes used to test batches of 16 PMTs for the PROSPECT detector.

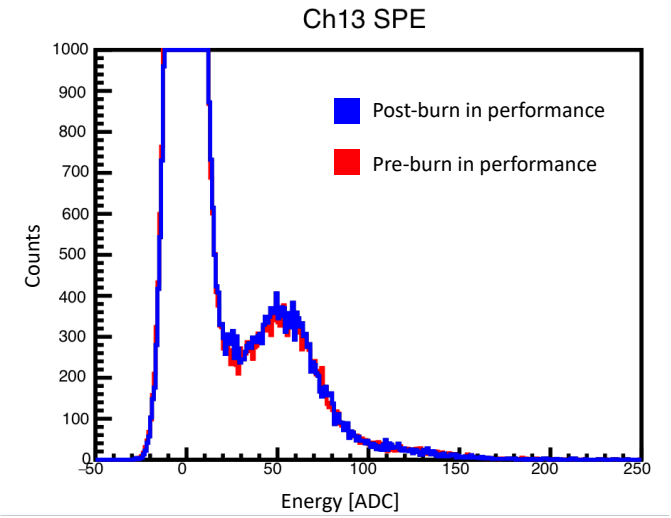

Figure 2.13: Example distribution of measured signals from pulsed LED source on a PMT with 1600V applied to it. The red distribution shows the measurement before burn-in, while the blue distribution is after burn-in. The strong agreement shows the PMT is stable over the 48 hour period, and qualifies for use in the PROSPECT detector.

measure the gain curve. The gain G is calculated from the fitted SPE mean in ADC as follows:

$$G = \text{ADC} \times 4\text{ns} \times 2\text{V}/2^{14}/50\Omega/q_e \tag{2.4}$$

where the DAQ utilized a 4ns sampling time, with a range over 2V with 14 bits, a 50Ω terminator, and q_e is the charge of an electron. An example of the gain curve for Hamamatsu PMT 68 is shown in Figure 2.14. The log of this gain curve was fit with a line, with the fit parameters stored in another SQLite database used in detector operation to ensure all PMTs were operating at approximately equal gains. While these coarse gain parameters were used throughout the entire detector operations, fine calibrations regarding energy reconstruction were done regularly and are described in greater detail in Sec. 3.3.4.

Any PMTs that failed to reproduce their SPE performance or had abnormal gain curves were re-evaluated and checked for faulty solder joints. After inspection, if no issues could be found in the joints, or if after fixing loose connections the PMT still exhibited abnormal behavior, the PMT would be designated as FAILED and would not be used for the

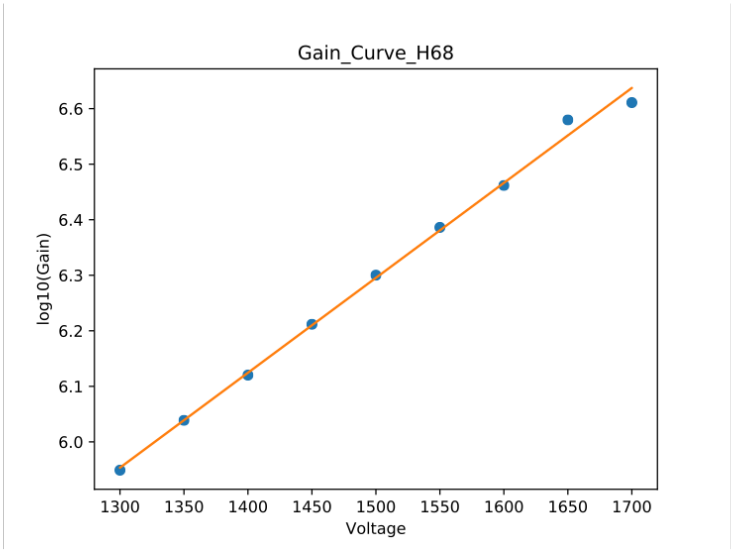

Figure 2.14: Example gain curve measured for a representative PMT for the PROSPECT detector. Gains were calculated using the fitted mean of the SPE peak as a function of applied voltage. A fit of this gain curve was calculated for each PMT and saved for use in operation of the full detector.

PROSPECT detector. Out of 256 (101) inspected Hamamatsu (ET) PMTs, 3 (16) did not pass qualification. Passed PMTs in housings were later filled with mineral oil, cleaned, and stored until detector assembly.

Final resistance checks of the cables, as well as mechanical preparation was done the day before individual housings were needed in assembly. All work on production housing assembly, characterization, and cleaning was done in a class 3000 cleanroom at Wright Laboratory.

2.2.5 Acrylic Vessel and Supports

The inner detector was assembled on top of the base of an acrylic tank using machined acrylic segments as support. These segments held the back plug of the PMT housings at the required 5.5° tilt and 0.146m pitch [54]. These pieces were assembled ship-lap style across the base plate as well as up along the sides of the detector in parallel with the assembly of the rows of PMTs. Additionally, vertical acrylic supports were mounted

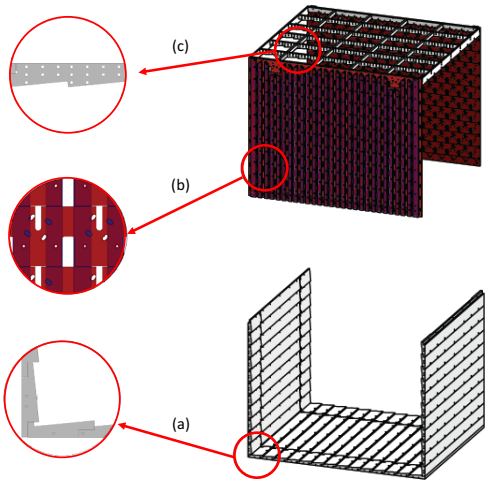

Figure 2.15: Acrylic segment support structure for the PROSPECT detector. (a) Wedge shaped panels supporting the PMT housing back plugs and support rods along the bottom and sides. (b) Vertical supports mounted to along the PMT end plates which support the housings and help rout cables. (c) Top baffles which strengthen the structure, hold top reflective panels in place, and allow a layer of LiLS on top of the inner detector volume [54].

along the backplugs of the housings. At the top of the support structure, machined acrylic baffles were attached for structural support and to hold the top layer of reflective separators in place. A diagram of the assembled bottom and side walls of the supports are shown in panel Figure 2.15a, and the interface of the vertical supports and top baffles is included in parts b and c of the same figure.

After assembly of the main detector volume, the side walls of the acrylic tank were lowered around the main detector volume onto the base and coupled via a double Viton O-ring seal and a tongue-and-groove joint. A strip of FEP was laid along the top of the side walls to cushion the third part of the acrylic tank, the lid. The lid, with fourteen rectangular holes ($5.1 \times 7.6 \text{cm}^2$) for cable and calibration tube feedthroughs, was lowered onto the side walls. All together the structure of the inner acrylic tank (inner dimensions $2.143 \times 1.995 \times 1.555 \text{m}^3$) as well as the acrylic support structures constrained the components of the optical volume of the inner detector. Photographs of the lid being lowered

Figure 2.16: (left) Photograph of the acrylic lid being lowered onto blocks on top of the acrylic sidewalls. (right) After blocking up the lid, cables and PTFE calibration tubes were fed through the holes in the lid. Then the blocks were removed, and the lid was lowered down onto an FEP strip along the rim of the side walls.

down onto the detector as well as the cables being dressed through the lid are given in Figure 2.16. After assembly, the entire vessel was bound together with a series sixteen of steel tension cables to secure the O-ring seals around the base of the tank. These were cushioned by 2.5mm-thick aluminum angles with 6.35mm-thick plastic strips and tensioned to 1300N with turnbuckles. After assembly, the gap between the O-rings was tested and held at a pressure of 7kPa via a small passageway installed near the base of the assembly. Additionally, a final seal of 5cm wide marine tape was applied around the outer seam as a back up.

2.2.6 Aluminum Container and Shielding

The inner vessel was placed inside an outer, secondary vessel made of aluminum which also protected the inner acrylic vessel during transport. Between the aluminum and acrylic

Figure 2.17: (left) Photograph of the interior polyethylene shielding located between the acrylic inner vessel and the aluminum outer vessel. (right) Photograph of cables being fed through holes in the lid of the aluminum vessel. Icotek fittings were later installed to seal the lid of the detector.

is a layer of 5% borated polyethylene (poly) for neutron shielding. A grid of poly was installed along the bottom of the detector so that the tension cables could fit undisturbed, and more poly pieces were shimmed into the volume after the inner vessel was in place. A picture of this inner poly shielding can be seen in the left panel of Figure 2.17. A layer of poly shielding was placed on top of the acrylic vessel before the cables and calibration tubes were fed through icotek fittings [67] in the lid. The lid was then closed over top of the detector, and another layer of poly shielding was placed on top

Once relocated to the HFIR site, additional shielding was installed on the detector, for a shielding package based on previous background measurements of the local backgrounds at HFIR [68]. An outer layer of high density polyethylene shielding was assembled log cabin style around a layer of 2.5cm thick layer of lead bricks, which were installed to mitigate ambient gamma backgrounds. A fireproof aluminum skin was attached around the detector, and layers of water brick shielding was installed on top of the detector for further neutron shielding along with a fiberglass fire blanket. Images of some of these shielding features can be seen in Figure 2.18.

Figure 2.18: (left) High density polyethylene shielding assembled in a log cabin style around the detector. (right) layer of lead bricks situated on top of the detector.

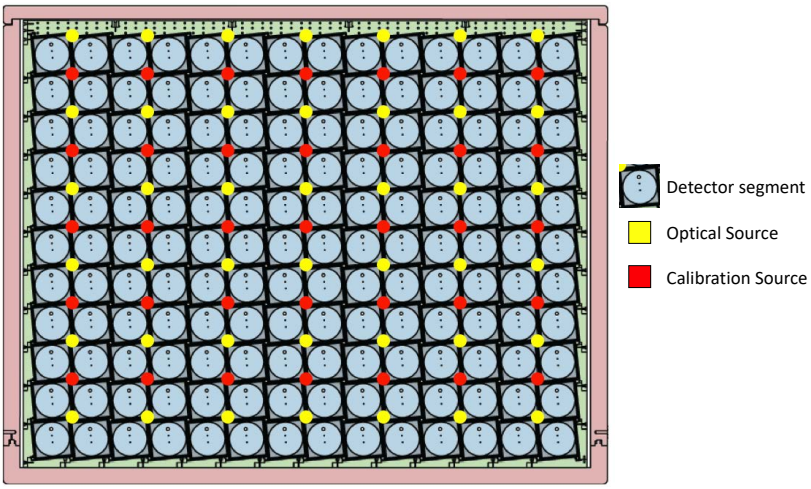

Figure 2.19: Locations of the radioactive source calibration tubes (red) and optical calibration sources (yellow).

2.2.7 Radioactive and Optical Calibration

As mentioned before, one of the main benefits of constructing the optical lattice on a slight pitch was to allow for the deployment of calibration subsystems along the axes of some vertices of the lattice. These separate locations are used both for deployments of the radioactive source calibration system as well as the optical calibration system. The specific locations for each of these deployments is given in Figure 2.19. The specifics of each system are detailed in the following subsections.

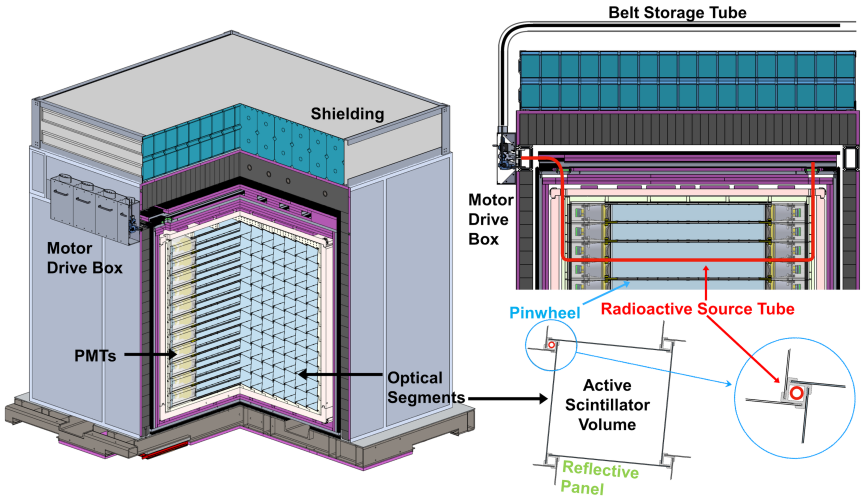

Figure 2.20: Interface of the radioactive source system with the detector. The left shows the motor drive boxes on the outside of the detector package. The top right gives a side view of an example path (red) of one of the source tubes through the detector. The bottom right shows the interface of the tubes with the optical lattice [69].

Radioactive Source Calibration

The radioactive source calibration system was designed to service the entire fiducial volume of the detector in a way that was chemically compatible with the LiLS, minimize dead mass inside the detector, and to operate within the tight reactor site constraints with remote operation and monitoring. This was implemented by deploying small, encapsulated sources along PTFE tubes via a single-ended length of timing belt using a system of stepper motors. The motor drives were mounted near the top of the exterior of the aluminum tank at the opening of a PTFE tube. The tubes went down to a specific location in the lattice of the detector, along the length of the optical segment through one of the PLA rods, and then back up to the top of the detector above the level of the LiLS. A sketch of an example path is given in Figure 2.20 in red, as well as its interface with one of the optical segment support rods [69].

The PTFE tubing was received on a roll and needed to be cut to length and straightened to reduce the stress put on the segment holders. This was done by annealing the tubes in

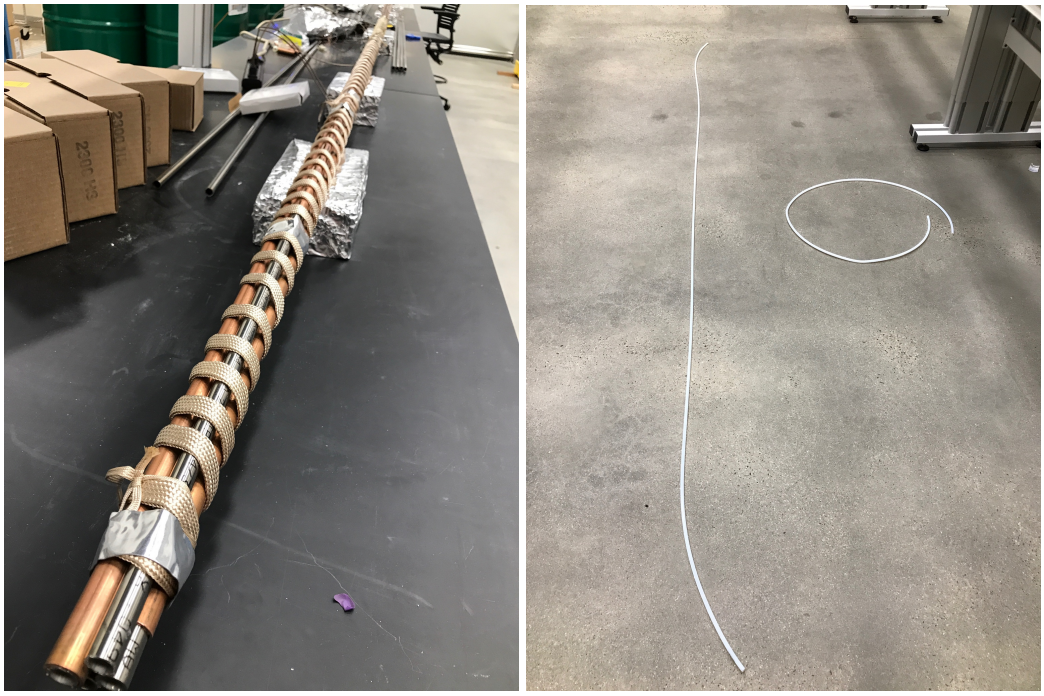

Figure 2.21: (left) Setup of the Oven for the Annealing of Teflon (OAT) used to prepare the PTFE tubes for the PROSPECT source calibration system. Heater tape, operated by a controller with two K-type thermocouples mounted inside the insulation (not shown), was wrapped around tubes of copper and stainless steel with an inner diameter matched to the outer diameter of the PTFE tubing. (right) Photograph showing the different behavior of the initial curved tubing and the straightened, annealed tubing.

a 20ft long linear oven consisting of four lengths of stainless steel tubes inside of copper tubes, all wrapped in heater tape, aluminum foil, and insulation. This Oven for the Annealing of Teflon (OAT) was controlled by two K-type thermocouples, installed within the insulation, and operated at temperatures of approximately 400°F. Photographs showing the setup without insulation, as well as the pre- and post-annealing PTFE tubes is given in Figure 2.21

The source drives themselves were NEMA 23 stepper motors mounted to 3D printed pulley assemblies. These assemblies consisted of 33 teeth pulleys slid onto the D-shaped motor shaft; the timing belt was held in place by spring loaded jockeys with an adjusted tension so that the motors can skip in the presence of unusual resistance such as a kinked tube. Photographs of these mounts and some of its components are shown in Figure 2.22

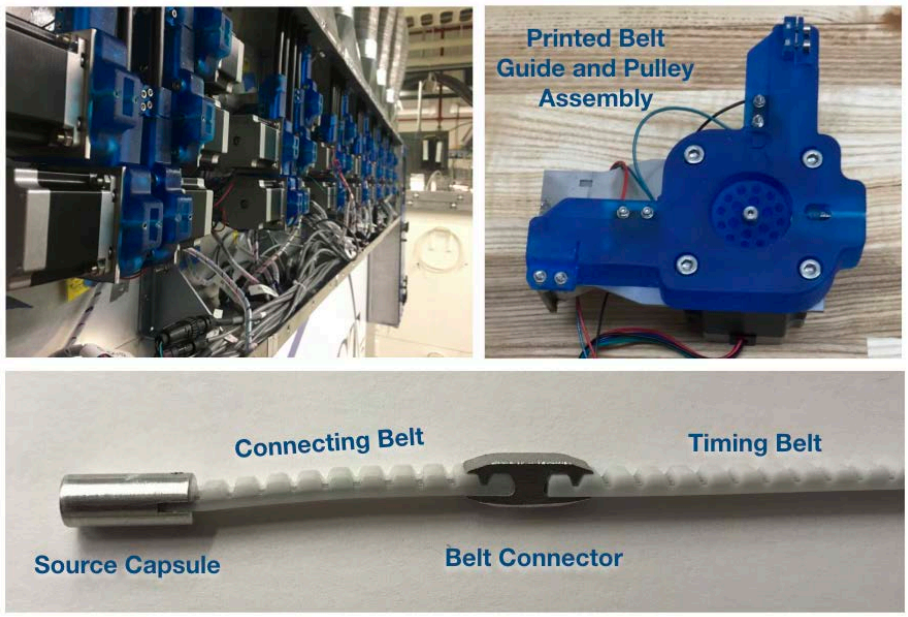

Figure 2.22: (left) System or source drivers mounted onto the PROSPECT Detector. (right) 3D printed assembly for the calibration timing belt and motor mount. (bottom) Exchangeable connecting belt with enclosed source capsule. It is attached to the end of a timing belt that is driven through the calibration tubes by a stepper motor [69].

[69]. An additional safety feature included was a microswitch that stopped the motor whenever the source retracted all the way up to the assembly, also acting as the homing mechanism for the motor drive.

The radioactive sources used for calibration consisted of enclosed capsules of several different sources for different purposes. β-decay sources were used to calibrate the full-detector energy scale as well as verify consistent segment-to-segment performance, e^+ sources were used to calibrate the effects of missing energy near detector edges, and spontaneous fission sources were used to calibrate the detector's neutron response. Specific sources used for calibration are included in Table 2.1 [55, 69].

In addition to energy reconstruction, sources were deployed at various positions along the longitudinal direction of the detector segments (z) to cross check position reconstruction, which is described in greater detail in Section 3.3.3.

Source	Decay Mode	Decay Energy [MeV]	Primary Purpose	Rate
^{137}Cs	beta-	0.662	Segment Comparison	0.1 μCi
^{22}Na	beta+	0.511, 1.275	Positron Reconstruction	0.1 μCi
^{60}Co	beta-	1.173, 1.332	Energy Scale	0.1 μCi
^{252}Cf	n(fission)	2.223 (n-H capture)	Neutron Response	866 n/s

Table 2.1: Radioactive sources deployed into the PROSPECT detector.

Optical Calibration

In addition to the radioactive source calibration system used for energy response, an optical calibration system (OCS) was implemented to help calibrate timing differences between PMTs in the same segment and monitor optical performance over time. The optical calibration starts with light produced by a 15 mW single mode fiber-pigtailed laser powered by a $<$10ns pulsed laser diode driver. The light, with a center wavelength of 450nm, is split into 42 separate light guides and piped through the calibration ports throughout the detector as shown in Figure 2.19. These light guides are optical fibers in 10 gauge Teflon sheaths for LiLS compatibility, which pass through icotek feedthroughs into the detector volume. At the end of the optical fiber, in the longitudinal center of the calibration location, is an optical diffusor. This diffusor consists of a small, conical reflector that distributes the light radially to all four adjacent segments through a Teflon diffusing cap. This cap also centers the diffusor in the segment via holes machined into the PLA support rod centers. These components are outlined in Figure 2.23 [55].

Besides the main 42 terminations of the OCS in the detector, 2 other splits of the original light go to two photodiodes to monitor the intensity of the laser output. Ultimately, this measurement showed that the laser splitter did not operate consistently in how it split the light from run to run, so relative monitoring via the OCS was not possible. Its position in the fixed center of the detector was still able to be used for position reconstruction cross checks, and as a source of light for SPE calibration for dedicated calibration runs to verify PMT consistency.

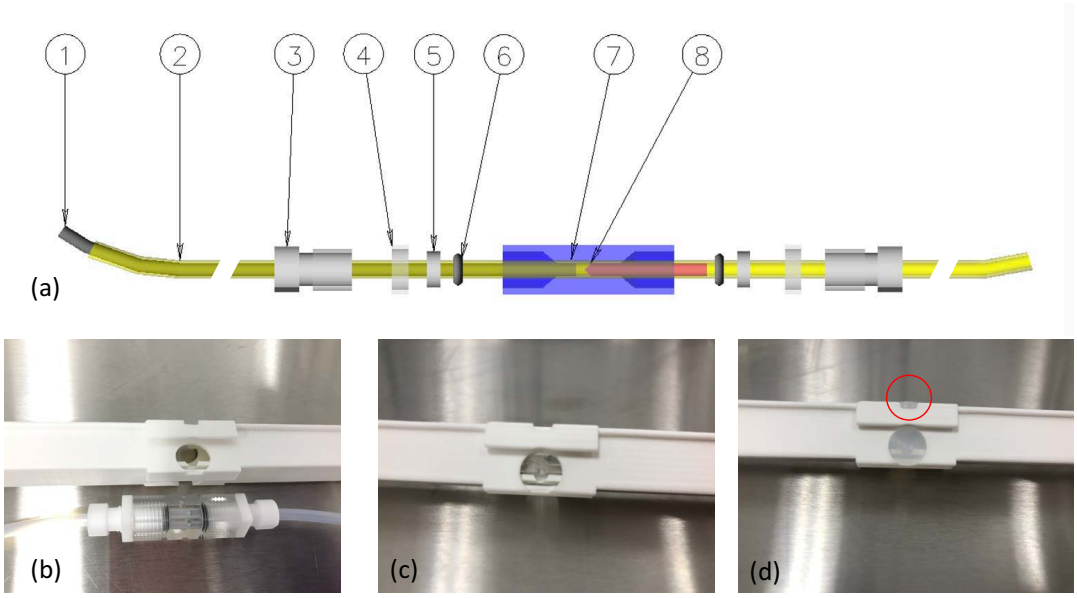

Figure 2.23: (a) exploded components at the terminating end of the optical source path: (1) fiber optic cable, (2) PTFE sheath, (3) compression nut, (4,5) spacer washers, (6) Viton O-ring, (7) square acrylic body, (8) conical reflector. (b) Photograph of the fully assembled optical assembly next to the PLA segment support, (c) optical assembly centered inside of the segment support, (d) completed assembly including Teflon diffusor caps. The edge face circled in red is the only part that is visible from the detector optical segment [55].

Parameter	Value
Active LiLS Volume	3760 L
Active target Mass	3680 kg ($/sim$4 tons)
Segmentation	11×14
Reconstructed z-position resolution	5cm
Center-to-center baseline (reactor to detector)	7.93 ± 0.1m
Baseline coverage	2m
Energy Resolution (RMS) at 1 MeV	4.5%
Signal-to-Cosmogenic Background Ratio	1.4
Signal-to-Accidental Background Ratio	1.8
Non-LiLS mass in the target region	3.4%

Table 2.2: Achieved parameters of the PROSPECT detector for its full data run.

2.3 Achieved Detector Parameters

Achieved detector parameters over the full dataset are summarized in Table 2.2 [55]. Baseline measurements done by a survey of various points in the reactor facility to points on the detector. These parameters all either meet or exceed the original design specifications of the PROSPECT detector, except for the baseline coverage which was limited by logistical constraints on moving the detector during operations. The PROSPECT physics reach was limited at smaller mass splittings because of this, but PROSPECT still explored new phase space at higher mass splittings without being able to move. Further details of the analysis (particularly with regards to the signal-to-background ratios) and results are detailed in the next sections.

Chapter 3

PROSPECT Event Reconstruction

One of the main goals of the PROSPECT analysis is to determine meaningful physics information from the raw digitizer output from the detector data acquisition system (DAQ). This conversion comprises of several stages and utilizes multiple types of calibration to fully reconstruct energy and position of events and clusters within the detector. This section will describe the DAQ system and the full process of converting from raw digitizer output to meaningful physics values. Beginning with the raw digitizer output, those data are unpacked into DAQ pulses. These pulses are then "crunched" into data frameworks called DetPulses containing information about the pulses coming from each PMT. To these pulses are applied calibration values, and the PMT-level pulses are then converted into segment-level physical values, which can then be clustered by time across multiple segments and used in the higher level physics analyses. A high level flow diagram of this is given in Figure 3.1, and each of the processes is detailed in the following subsections.

3.1 Unpacking DAQ Output

The PROSPECT detector comprises of 308 individual PMTs that feed into channels of twenty-one waveform digitizers (WFD) (CAEN-V1725 [70]). These WFDs have a 250MHz

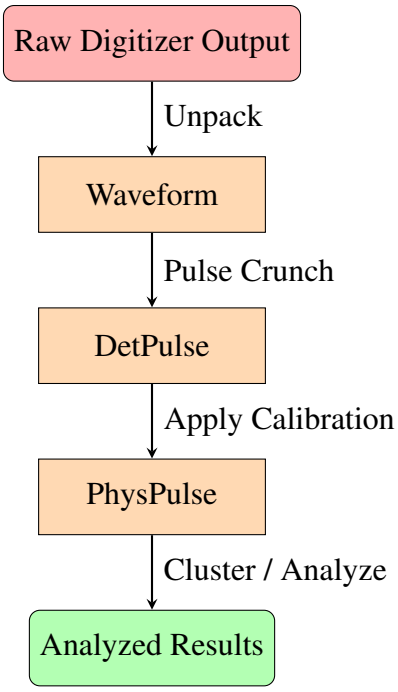

Figure 3.1: Schematic of analysis workflow for the PROSPECT data analysis. Orange rectangles represent various intermediate steps of the analysis, while the arrow labels refer to the processes being performed to reach the next level of the analysis.

sampling rate and a 14 bit depth per sample and are operated in two VME crates [71] with synchronized clocks and a hardware level triggering system. A schematic of the DAQ layout is given in Figure 3.2 [54].

The WFDs are all connected through a custom Logic Fan-In/Fan-Out PS-FIFO [72], so that all PMTs can trigger off of a coincident event between PMTs in one segment, each above a signal of approximately five photoelectrons within 64ns of each other. After the initial trigger, all channels above a smaller threshold are recorded to disk through use of a zero-length encoding (ZLE) threshold. Figure 3.3 [55] shows an example where the two top-most waveforms surpass the trigger threshold to trigger data collection, while all waveforms that surpass the ZLE threshold are recorded to disk. The ZLE is crucial for maintaining reasonable data rates throughout detector operation, which would be limited by the 85MB/s optical fiber bandwidth without it. During reactor-on operation, the total acquisition event rate is $\sim 20 \cdot 10^4 \mathrm{s}^{-1}$, which with the ZLE is approximately 3MB/s. These

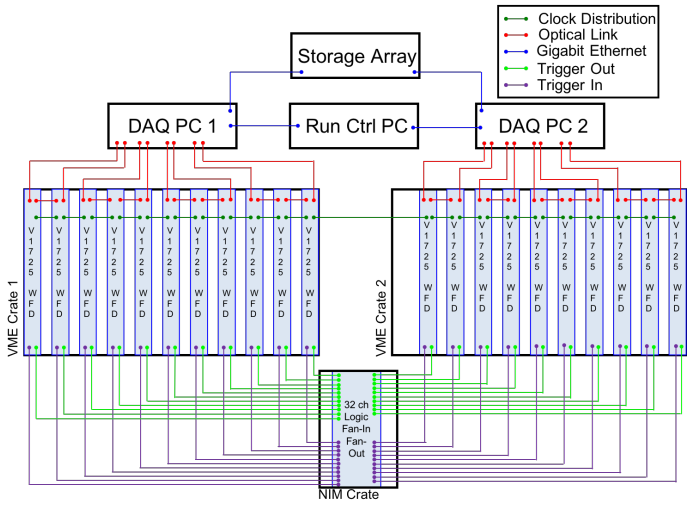

Figure 3.2: Schematic of the PROSPECT data acquisition system [54].

waveforms are encoded in binary on one of two buffers on the WFD modules, one which is being filled with new data and one which is available for transfer to disk. This data is transferred to the DAQ control PCs and written to disk as a compressed binary file for each digitizer board per one hour run.

Offline, these data are unpacked into TTrees in ROOT [73] files containing all of the waveforms for each channel into a single file for each ∼1h run. This unpacking process includes a logical channel map which maps each hardware channel based on headers in the binary files to logical functions (e.g. specific PMTs in the detector). The main attributes of each TTree are the event numbers, PMT channels, and the arrays of waveform sample values. XML files containing metadata of the runs are produced and updated throughout the data analysis framework through use of an SQLite3 [74] runs database.

3.2 Crunching Waveforms

Once stored as proper waveforms, data files were processed with the custom PulseCruncher program. This program calculates a number of meaningful values from each waveform. After identifying peaks as any local maxima (separated from other peaks by at least 20ns),

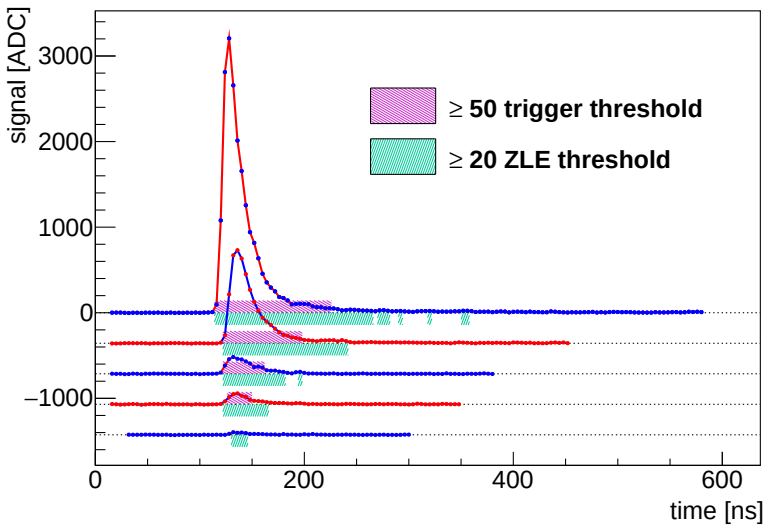

Figure 3.3: Example waveforms of multiple channels showcasing the PROSPECT triggering procedure. All waveforms are inverted and include a y-axis offset for visual clarity. Blue (red) waveforms correspond to PMTs on the west (east) side of the detector. The magenta regions correspond to signals that surpass the 50ADC trigger threshold while cyan regions correspond to regions that surpass the 20 ADC ZLE threshold and will be recorded to disk [55].

the time of half the max height is used to reference all time windows, calculated as the linear interpolation to 50% of the max height of the pulse. The baseline is set as the median of all values 5 to 30 samples before the leading edge, and is subtracted off from each value in the waveform. The area of the waveform is integrated from 3 samples before the leading edge to 25 samples after the leading edge. Finally, a PSD metric is defined as the integrated ADC in the tail of the pulse divided by the total integrated ADC of the pulse, where the tail is integrated from 11 samples after the leading edge to 50 samples after the leading edge, while the full integral for the PSD metric is defined over the range of 3 samples before the leading edge to 50 samples after the leading edge. Figure 3.4 [54] shows the different regions used to calculate these values. The final output of this is a DetPulse ROOT file for each $\sim$1h data run with a TTree containing the event numbers from the previous step, the PMT channel of each event, the arrival time from the start of the file, the pulse area, the max height, the baseline, the leading edge rise time, and the PSD metric.

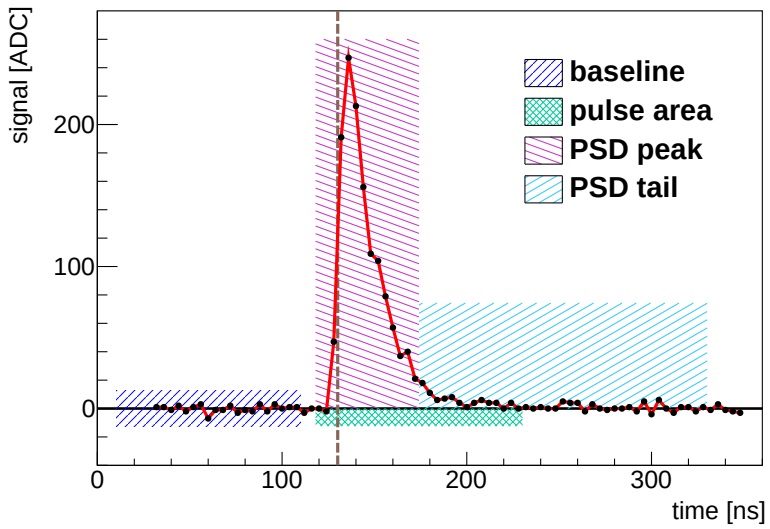

Figure 3.4: Representative regions used for Crunching a waveform. The location of the half height of the leading edge is used to define the timing. The regions used to calculate the baseline, pulse area, and PSD metrics are shown [54].

3.3 Calibrating Detector Pulses

The next step in data processing involves several steps of calibration. These processes are managed with a SQLite3 calibration database setup similarly to the runs database. Entries are managed on a segment-by-segment manner, and all of the following events require DetPulses to be coincident between both PMTs in a segment (within 20ns) for any event.

3.3.1 Timing Calibration

Event timing in a segment t_i is calculated as the average of the timestamps of the times seen by each PMT: $t_i = \frac{1}{2}(t_0^i + t_1^i)$. Similarly, a metric is defined as the difference of the individual PMT times $\Delta t^i = t_0^i - t_1^i$. These metrics are calibrated using through-going cosmogenic muon tracks in the detector, with muon events selected as any events with an ADC sum $> 10^5$ coincident in at least 4 paired segments. While muons that pass through the full width of a segment deposit so much energy that the signal exceeds the dynamic range of the digitizer and saturates the signal, muon tracks that clip the corner of

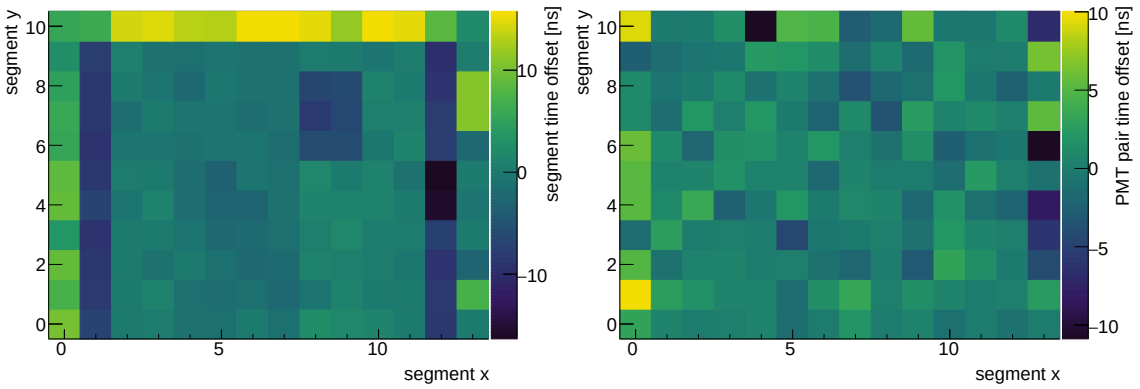

Figure 3.5: (left). Example 2D histogram from one calibration period showing inter-segment timing offsets $\overline{t}^{i}$. Notable features such as the larger deviations from zero around the sides and top row of the figure arise from the different spread in performance between the ET and Hamamatsu PMTs as well the clusters of offsets like on x segments 1 and 12 are due to timing offsets between boards (right) Example showing intra-segment PMT timing offsets $\overline{\Delta t}^{i}$ [54].

the segment deposit sufficiently low enough energy to produce well behaved signals that can be used for calibration.

Since muon tracks provide simultaneous signals in the detector, up to the muon transit time between segments, the total segment-to-segment timing offsets are determined by averaging the values $T^{ij} = t^{i} - t^{j} - t^{ij}_{\mu}$ where t^{ij}_{μ} is the muon transit time between segments i and j as calculated by a Principal Components Analysis (PCA) trajectory fit to pulse pair data. Solving a linear system with all values of $\overline{T}^{ij}$ provides offsets for each segment up to a global offset, and similarly $\overline{\Delta t}^{i}$ can be used to calibrate the relative offsets of PMTs within a segment. Muon statistics are large enough to be used for calibration on a run-by-run basis. A plot of the both the segment-level offsets and the PMT-level offsets are shown in the two panels of Figure 3.5 [54] from one calibration period.

3.3.2 Combined PSD

The PSD parameters from each PMT in a segment must be combined to establish a single PSD value for the event, and positional dependence of the individual PSD signals must

be accounted for. This positional dependence is explained by a generally wider spread in photon transit distances to the photocathode as the initial event is farther from the individual PMT. In order to account for this positional dependence, the PSD of corner clipping muon events for each PMT is plotted as a function of Δt and fit to the three-parameter curve PSD$(\Delta t) = p \cdot (1 + d \cdot [1 - e^{k \cdot \Delta t}])$. Here p corresponds to average PSD at the center of the segment, d represents the depth of the effective PSD drop off as a function of Δt, and k represents the effective variation length scale of the drop off. These fitted curves are then subtracted off of the real PSD values in order to remove the positional dependence on the event, leaving electron-like events centered around p. This new, position independent PSD metric is averaged between the two PMT signals for a given segment event weighted by the estimated number of photoelectrons in each pulse. In order to define PSD selection criteria for event selection later down the analysis chain, the fitted value and width of p is tracked, and selection of electron-like events is defined in terms of significant deviations from p.

3.3.3 Position Reconstruction

Longitudinal position is calibrated for events using both the timing distribution between PMTs in a segment Δt as well as the log ratio of the pulse areas seen by each PMT $R \equiv \ln \frac{S_1}{S_0}$. A distribution of Δt from corner-clipping muon events is used to set a preliminary position z. While internal features are visible, the sharp drop offs in the distribution at the ends of the cell are sufficient to establish the position in the segment to $\sim$1cm. The distribution of Δt is fit to a linear polynomial plus a cubic term (the quadratic term is omitted assuming an antisymmetric correction factor). The correction factor primarily accounts for edge effects where the optical path becomes more complicated near the front faces of the PMT optical modules. Figure 3.6 [55] shows these event rates vs both the Δt value and the geometric mean of the pulse areas $\sqrt{S_0 S_1}$, which is a mostly position

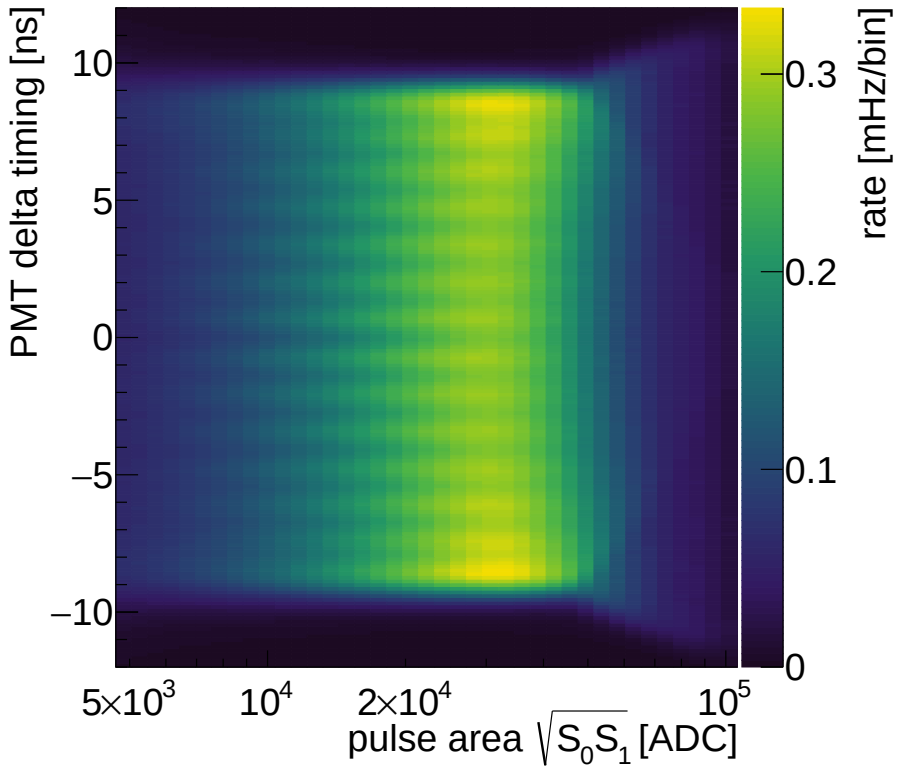

Figure 3.6: 2D histogram showing event rate vs PMT Δt and pulse area. The horizontal stripes are a result of the Hobbes effect, where the diffuse tabs from the segment holders cause perturbations in light transport at regular intervals. Event rates are fit as a cubic function of Δt using the sharp drop offs to mark the ends of the segment to partially determine reconstructed z-position. The tiger stripes from the Hobbes effect are used to cross check the distributions [55].

independent proxy for the event energy. Additionally the internal features in the distribution, which result from light transport perturbations due to the diffuse components of the segment holders in the optical volume, produce horizontal stripes in this distribution. Since these features have well known positions from detector construction, these features from the Hobbes effect are used as a crosscheck for the position reconstruction.

Using this estimate of z from Δt, R can also be fit to a linear function plus cubic term to determine another estimate for the reconstructed position in a similar manner. The numerical terms for the fits for both $z(\Delta t)$ and $z(R)$ are stored in the calibration database, and a final reconstructed z-position z_{rec} for each event based on a statistically

weighted combination of the two estimates. Additionally, z_{rec} is cross checked using a comparison of measured position to nominal locations of radioactive sources deployed using the calibration system. The absolute precision of the nominal positioning of sources is limited to a few centimeters, and is not more accurate than using cosmogenic muons for calibration.

The total uncertainty on z_{rec} is measured using various coincident α-decays from isotopes of Po present in the detector from an added trace amount of ^{227}Ac ($\sim$3 s^{-1} α-decay rate) and natural radioactive backgrounds from the decay chains of ^{238}U and ^{232}Th. Distributions of event rates vs z_{rec} are shown in the left panel of 3.7 [55] for each of these isotopes in a single segment of the detector. These events are selected using coincident time cuts, PSD, and reconstructed energy (which is described in the next subsection) as outlined in Table 3.1 [55].

Decay	Selection Criteria				Rate
	E_{rec} (MeV)	PSD	Pulses	Δt_{rec} (μs)	(mHz)
^{219}Rn α	(0.57,1.15)	(0.19,0.36)	1	(0,5000)	403
^{215}Po α	(0.66,1.15)	(0.19,0.36)	1		
^{214}Bi $\beta+\gamma$	<4.00	(0.05-0.22)	Any	(10,710)	165
^{214}Po α	(0.72,1.00)	(0.17,0.34)	1		
^{212}Bi $\beta+\gamma$	<3.00	(0.05-0.22)	Any	(0.7,1.7)	55
^{212}Po α	(0.95,1.27)	(0.17,0.34)	1		

Table 3.1: Selection criteria and rates for correlated decay signals in PROSPECT used for performance evaluations. For bismuth decays, given PSD cut values are applied to the highest energy pulse in the cluster; relaxed time-dependent PSD cuts are also applied to other pulse clusters. Integrated rates include only segments used in the oscillation and spectrum analyses.

These distributions are given in the left panel of Figure 3.7 [55]. The sharp drop offs at the edges of the distributions, as well as the distribution of the difference in reconstructed position of the localized α-decays for the ^{215}Po decays, are consistent with a z_{rec} uncertainty of $\sim$5cm.

These techniques are all used to calibrate z_{rec}, but due to the relatively low event rates

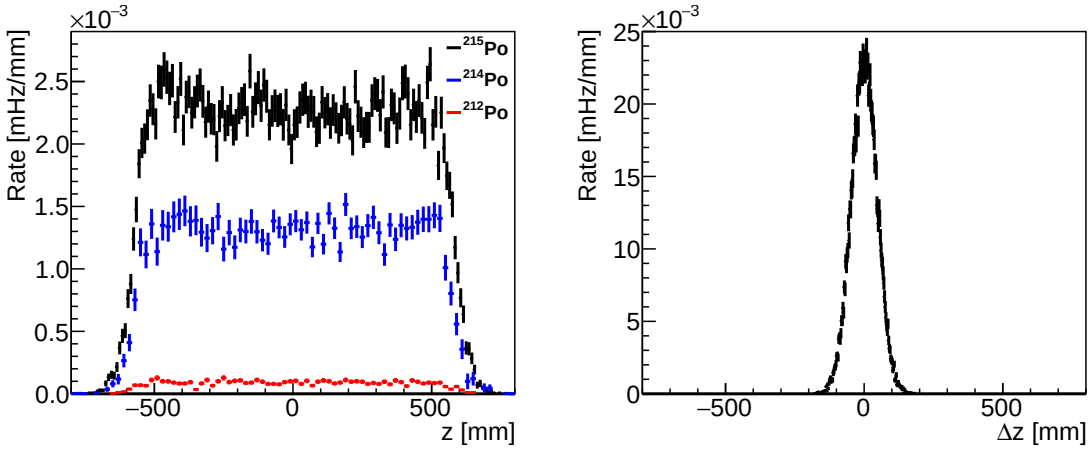

Figure 3.7: (Left) Reconstructed positions of the α-decays for backgrounds distributed uniformly across the length of the PROSPECT detector. (Right) distribution of the difference in z_{rec} between the two, localized coincident events ending in the α-decay of ^{215}Po. The width of this distribution corresponds to the uncertainty on z_{rec} of $\sim$5cm.

they must use statistics on the week-long time frame. Thus the whole dataset is subdivided into 11 calibration periods for position reconstruction purposes.

3.3.4 Energy Calibration

The segmentation of PROSPECT, as well as scintillator quenching and trigger acquisition thresholds, leads to a complicated energy response. Rather than completely reconstructing all energies to the energy of the original IBD on an event-by-event basis, energy response parameters are incorporated into a detector response function for the full detector. This process is done in two stages. First, the visible energy E_{vis} is a measure of the energy seen by the detector, which incorporates effects like scintillator quenching (which limits the light seen for detected events) but calibrates out positional- and time-dependent effects of the detector. After E_{vis} is constructed, an energy response model is fitted using a variety of energy calibration techniques to properly tune a complete simulated model of the PROSPECT detector response. This detector response matrix is then a key aspect of the later analysis and incorporates effects from energy deposition in inactive volume of the detector, scintillator non-linearity, and data acquisition thresholds.

The visible energy E_{vis} is defined by a weighted sum of the pulse areas from the two PMTs in a given segment S_0 and S_1 as

$$E_{vis} = \frac{S_0 n_0/g_0 + S_1 n_1/g_1}{n_0\eta_0(z_{rec}) + n_1\eta_1(z_{rec})} \tag{3.1}$$

where n_i/g_i is the expected ratio of counted photoelectrons to pulse area at the segment center for PMT i, and $n_i\eta_i$ is a position-dependent light transport efficiency correction factor ($n_i\eta_i$ is normalized to the expected number of photoelectrons / MeV at the segment center). For later parts of the calibration, the energy scale is fixed by the neutron capture on ^{6}Li, which provides a continuous, monoenergetic signal throughout data taking that is easily distinguishable above γ-ray backgrounds by its PSD metric. The neutron capture event is scaled to fall at E_{vis} = 0.526 MeV. The neutron capture event is also used to track the gain parameter g_i as well as the light transport curves η_i for each PMT as a function of time. In particular, the photostatistics per MeV n_i is determined by the width of neutron capture peak for each PMT for each calibration period, and η_i is similarly fit as a function of z_{rec}. Gain calibration constants g_i are calculated on a run-by-run basis while roughly two-week time periods are used as calibration time windows to have sufficient statistics for n_i and η_i calibration constants.

Degradation of the scintillator optical properties over the entire data run for PROSPECT led to a decreased light collection per MeV, which led to an overall deterioration in reconstructed energy resolution. This can be seen directly in the change in average light curves of PMTs from the beginning to the end of the seven month data taking period as shown in Figure 3.8 [55]. The red lines show the initial light curves $n_i\eta_i(z_{rec})$ from opposite PMTs averaged across all channels. The blue lines show the averaged light curves from the end of the data set, where the bands represent the RMS spread between channels. This drop in the collected light curves (approximately 50% after the summing done in Equation 3.1) is a result of both decreased light yield and effective attenuation length in the detector.

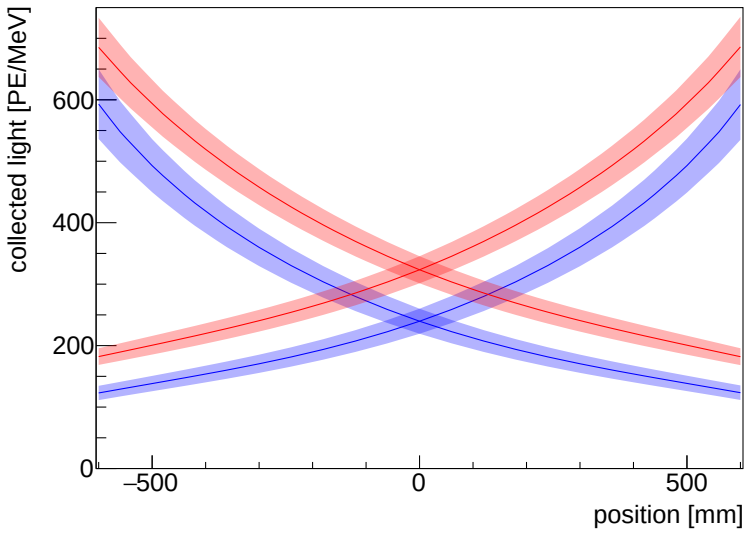

Figure 3.8: Average light collection curves for each opposite PMTs in the detector as a function of reconstructed event position along the optical segment. The upper red curves show the average collection at the beginning of the data set, while the lower blue curves show the light collection at the end of the data set. The total loss (from summed signals) is roughly 50% over the 7 month time period. Bands are calculated from the RMS spread between channels [55].

Because reactor-off data is subtracted from reactor-on data in order to directly account for correlated backgrounds, a uniform response matrix is needed across the entire data taking period. This means that if the deteriorating E_{vis} resolution would lead to improper background subtraction if left as is. In order to unify the analysis approach, a new smeared energy variable is introduced E_{smear}. By adding an artificial energy smearing to the earlier data commensurate to its difference from the final energy smearing, this simplification sacrifices the better energy resolution of earlier stages for a simpler, more direct analysis. Using a decay from the spiked ^{227}Ac source, an example of the gradually broadening resolution of E_{rec} is shown compared to the relatively stable $E_{smeared}$ resolution across the entire data taking period in Figure 3.9 [55]. While sometimes E_{rec} is used in the notation in this manuscript, all final analyses use the $E_{smeared}$ metric including the IBD selection and analysis as well as the joint analysis with Daya Bay.

Finally, to be clear of effects from variation in the relative energy scale of segments

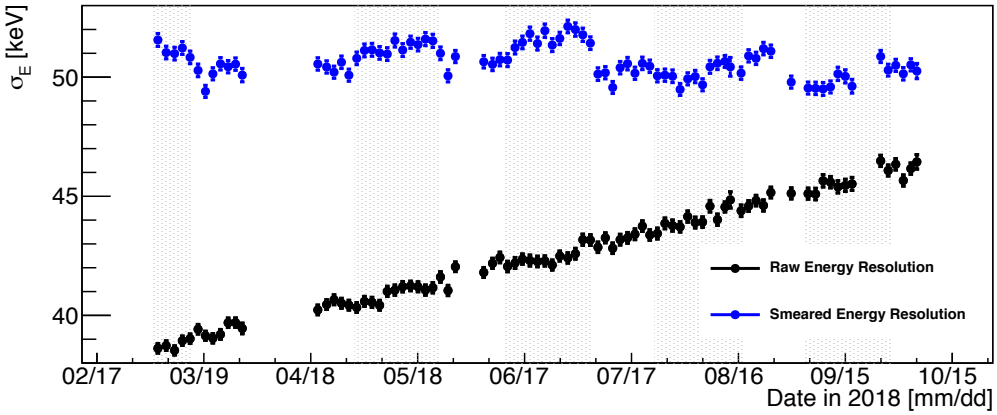

Figure 3.9: Resolution of E_{rec} at the ^{215}Po peak from ^{219}Rn-^{215}Po $\alpha - \alpha$ decays. Black shows the original E_{rec} with its broadening over the data taking period while blue shows the relatively constant E_{smear} variable, which has been artificially smeared for a consistent performance throughout the entire data set [55].

on their hardware level ZLE threshold described in Section 3.1, an analysis threshold is set at $E_{smear} > 90$keV. This threshold was determined from toy studies that found the minimum threshold that fully removed the effects from relative energy scale variations. This threshold slightly decreases the total E_{rec} of events, but ensures uniformity across the volume of the detector.

3.3.5 Segment Clustering

After the detected energy is calibrated at an individual segment level, events are clustered by time to obtain a more comprehensive event reconstruction. Grouping any segment-level events that are no more than 20ns apart from their nearest neighbor in time, events are clustered and given a timestamp based on the median time of individual energy depositions. The total energy of the cluster is the sum of each E_{smear} of events in the cluster, and the cluster segment and z_{rec} is assigned by the segment-level event with the greatest energy deposition. A single cluster may have a multiplicity greater than one where energy is detected in multiple segments, and all individual E_{rec}, PSD values, and z_{rec} are saved for those clusters. Later on in the higher level analysis, clusters can be defined to be in

coincidence with each other across larger time scales, such as IBD events with prompt and delayed clusters occurring within approximately 120 μs of each other. The cluster-level variables for the cluster E_{smear} and segment are the main inputs used in the high level oscillation and spectrum analyses.

3.4 Detector Energy Response

Once the visible energy is completely reconstructed and clustered for the data, the global detector energy response can be calculated using the robust, custom built simulation framework called PG4 based on the GEANT4 software package [75]. In PG4, the full detector geometry is simulated using as-measured dimensions for nearly all components of the detector including the scintillator volume itself as well as features of the optical grid (including support rods, materials, and dimensions), as well as radioactive source capsules and other material volumes including air, acrylic, Teflon, and PLA. The simulation also includes the geometries of the PMT housings, acrylic support structure, both acrylic and aluminum tanks, and the inner and outer shielding packages.

In the simulations, the Monte Carlo non-linear energy response E_{MC} of the scintillator light from true energy deposition is modeled using

$$E_{MC} = A \sum_i \big(E_{scint,i}(k_{B1}, k_{B2}) + E_{c,i}(k_C) \big). \tag{3.2}$$

where i are the individual steps in the simulation. In this equation, the total energy is based on contributions from scintillation light E_{scint} and Cherenkov light E_c. E_{scint} is defined by

$$\frac{dE_{scint}}{dx} = \frac{\frac{dE}{dx}}{1 + k_{B1}\frac{dE}{dx} + k_{B2}(\frac{dE}{dx})^2} \tag{3.3}$$

where $\frac{dE}{dx}$ is the true energy deposition in the volume and k_{B1} and k_{B2} are the first- and second-order parameters from Birks' law quenching [76]. E_c gives the energy from

Cherenkov light that is produced, absorbed, and then re-emitted through scintillation and is modeled at step i as

$$E_c = k_c \sum_{\lambda} N_\lambda E_\lambda \tag{3.4}$$

where N_λ is the number of Cherenkov photons emitted per unit wavelength, E_λ is the energy of those photons, and k_c is a relative scaling factor. The summation in Equation 3.2 is scaled by a total normalization A to allow for different fractional rates of the conversion of deposited energy into detected energy.

For a given event simulation, each step $E_{MC,i}$ is tracked and summed on a segment-by-segment basis, and the total detected energy is converted into a simulated PMT waveform based on templates maintained in PG4. The shape of the template is defined by the amount of total quenching in the event, while the amplitude of the simulated pulse is set by the magnitude of E_{MC} and a positionally-dependent factor. After the waveforms are generated for events, they are processed using the same analysis as real data.

In order to accurately describe the detector energy response, the four parameters A, k_{B1}, k_{B2}, and E_c from Equation 3.2 are fitted using calibration data. Data from deployments of gamma sources ^{60}Co, ^{137}Cs, and ^{22}Na, as well as n-H captures from a deployed ^{252}Cf neutron source, and cosmogenically-produced ^{12}B are used. The deployed sources have features in their energy spectra associated with each of the decays listed in Table 2.1, covering a ranged of features from 0.5MeV to 2.5MeV in energy, and all sources were deployed at the detector center in order to most efficiently capture energy in the detector. The ^{12}B events are selected from a prompt, nuclear recoil like event (selected by PSD cut, with $0.7\text{MeV} < E_{rec} < 10\text{MeV}$) followed by a delayed, electron-like event. The two events must occur within 3ms - 30ms and be low multiplicity events within 12cm of each other, which identifies ^{12}C(n,p)^{12}B events and the subsequent β-decay of ^{12}B. By using the endpoint of the ^{12}B β-decay as a feature, this can extend the calibrated range up to 12 MeV.

To fit the four energy scale parameters, the following χ^2 is constructed and minimized:

$$\chi^2_{data-MC} = \sum_{\gamma} \chi^2_{\gamma} + \sum_{multi} \chi^2_{multi} + \chi^2_{^{12}\mathrm{B}} \tag{3.5}$$

where $\sum_{\gamma} \chi^2_{\gamma}$ is calculated from the difference between data and MC for the three γ sources as well as the subsequent γ-decay after the n-H capture associated with the ^{252}Cf source, $\sum_{multi} \chi^2_{multi}$ is the difference between data and MC for the multiplicity distributions for those same sources, and $\chi^2_{^{12}\mathrm{B}}$ is calculated from the difference between data and MC for the ^{12}B β-spectrum. From the best fit values for the energy scale parameters, the MC is in excellent agreement by eye with the calibration data, as showcased in Figure 3.10 [55]. The best fit χ^2 / degrees of freedom is 581.5/420. In the final best fit, the contribution from Cherenkov emission was minimal relative to the direct scintillation light. For the n-H capture events, for example, E_c contributed only 3.5% to E_{MC}.

Energy scale calibrations were done using dedicated data runs in April, May, August, and December of 2018. The April dataset was used to determine the detector energy scale parameters for the full analysis, while the other data sets were used as crosschecks. The top panel in Figure 3.11 [55] shows the ratio of E_{rec} of data to MC for several features in the 0.5 - 2.5 MeV range, while the bottom panel also includes the corresponding features from the August and December calibration campaigns, including an extra feature from a deployed AmBe source in December. All features in the γ calibration ratio are found consistent within 1% for the main calibration set, and all features in the crosschecks are similarly within uncertainties of the of the prediction, showing stability of the non-linearity of the detector energy response.

Similarly from the energy scale, the energy resolution is also fit to calibration data. The relative energy-dependent resolution $\frac{\sigma_E}{E_{rec}}$ is defined by

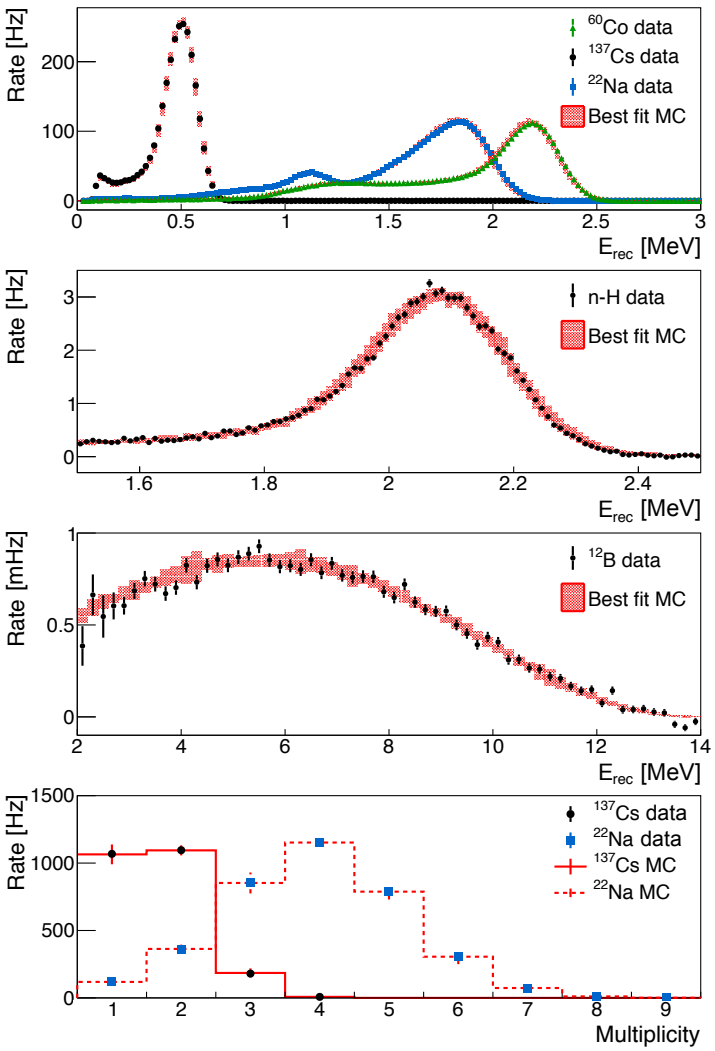

Figure 3.10: Comparison of data and MC used to determine detector energy scale. Simulated response is shown in red in each panel. Overall, there is excellent agreement between data and simulated distributions. (Top) Reconstructed energy spectra comparison from detector center-deployed γ sources. (Top Middle) Reconstructed energy spectrum from n-H captures from ^{252}Cf center deployed-source. (Bottom Middle) Reconstructed energy β-decay spectrum from cosmogenically produced ^{12}B. (Bottom) Event multiplicity distributions from representative sources. ^{137}Cs and ^{22}Na are the lowest and highest average event multiplicity respectively, but ^{60}Co and ^{252}Cf n-H capture multiplicity distributions were also used in determining the detector energy scale best fit parameters [55].

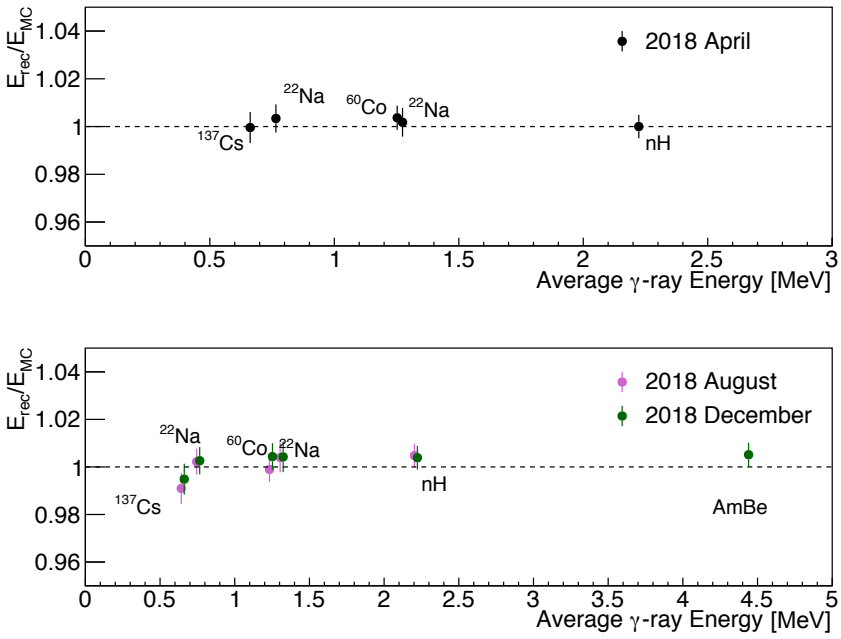

Figure 3.11: (Top) Ratios of reconstructed energy for calibration energy peaks in data over predictions from MC. Results from the April 2018 calibration run which were used to set the energy scale parameters for the analysis. (Bottom) Corresponding energy features from later calibration runs from August and December 2018 used to crosscheck the energy scale parameters. The December data includes an additional energy feature from a deployed AmBe source, which is also in excellent agreement with predictions from simulation. [55].

$$\frac{\sigma_E}{E_{rec}} = \sqrt{a^2 + \frac{b^2}{E_{rec}} + \frac{c^2}{E_{rec}^2}} \tag{3.6}$$

where a corresponds to light collection inefficiencies, b corresponds to energy-dependent photostatistics, and c corresponds to electronic noise. This equation is fit against fitted resolutions of spectral features in the calibration γ sources. The best fit energy resolution is included in the right panel of Figure 3.12 [55]. This demonstrates the intrinsic energy resolution of the PROSPECT detector at approximately 5.5% relative energy resolution at 1 MeV.

After setting the parameters of the detector energy response, a detector response matrix is generated using the full simulation and analysis framework. IBD interactions are

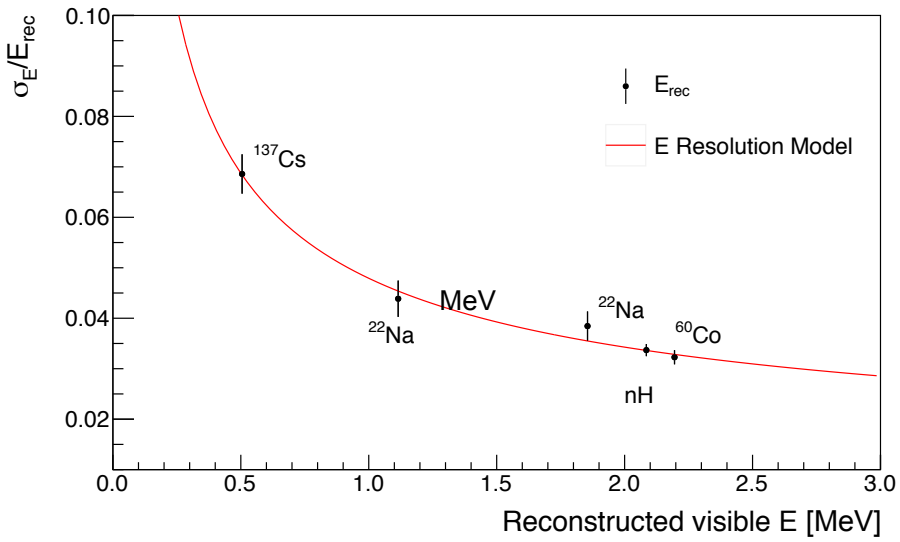

Figure 3.12: Relative energy resolution as a function of E_{rec}. Red shows the best fit Energy resolution function based on peaks from calibration data sets [55].

produced from neutrinos uniformly distributed in energy from 1.8 - 10 MeV. These events are used as inputs for the simulation waveform generators, passed through the entire IBD event selection analysis detailed in Chapter 4, and filled into a 2D histogram based on their original antineutrino energy and their final reconstructed energies. The histogram is generated using 50 keV wide bins along both axes, and is globally normalized so all elements sum to one, maintaining relative efficiency as a function of energy. This matrix is used both in the final analysis to produce predicted spectra as well as in studies on systematic uncertainties. The matrix itself is shown in Figure 3.13 [55] as well as a monoenergetic slice of the matrix to show the reconstructed energy spectrum of a 4 MeV antineutrino in Figure 3.14 [55]. A few notable features of the PROSPECT detector response are the effects from energy resolution (dominated by light collection), a downward shift in reconstructed energy from antineutrino energy driven primarily by quenching in the LiLS, and a peak at 511 keV in reconstructed energy (visualized as a horizontal line in the figure) associated with detected annihilation gammas from events with a primary vertex outside

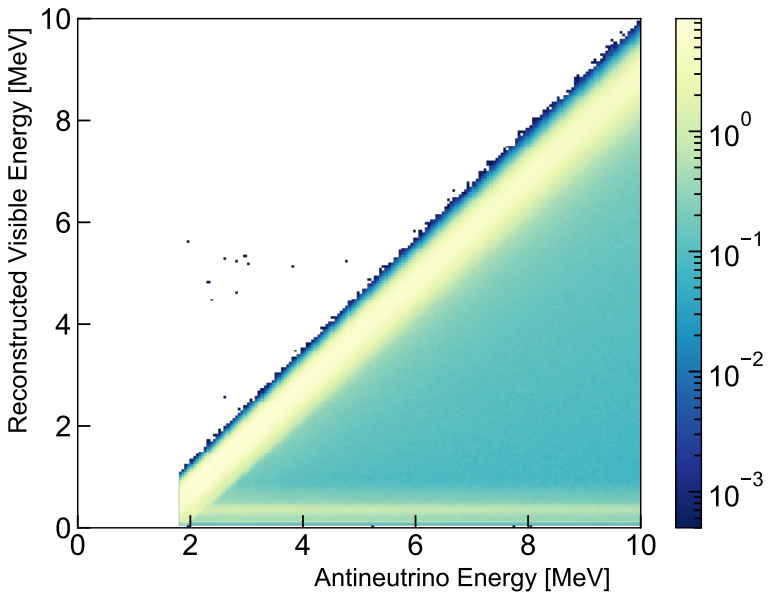

Figure 3.13: Full detector response matrix for PROSPECT. This matrix maps the transformation from a true antineutrino signal to the observed reconstructed visible energy measurement. [55].

of the active volume of the detector.

With the completion of the event-by-event calibration on the basis of position, energy, PSD, and a dedicated detector response function, the data is stored as a new PhysPulse object (as opposed to the DetPulse objects with the waveform relevant data) containing physically motivated parameters for each event. These parameters include the event number, index of the optical segment, reconstructed visible energy, time stamp, reconstructed position along z, and PSD value for each event in the detector in 1 hour data files. These values, as well as the values previously mentioned for the clustering procedures, are used in the subsequent IBD analyses, both for the sterile neutrino oscillation analysis and the ^{235}U spectrum measurement. The details of that event selection and higher level analyses are described in Section 4.

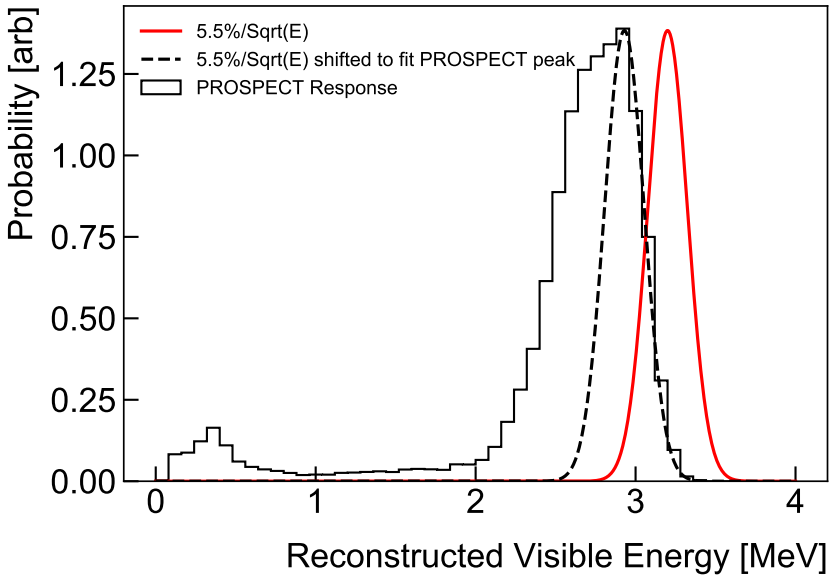

Figure 3.14: Monoenergetic slice of response matrix at 4.0 MeV antineutrino energy. The response function includes an energy resolution, an average downward shift in energy due primarily to quenching in the LiLS, and a peak in reconstructed energy around 0.5 MeV associated with captured annihilation gammas from events with primary vertices in inactive volume of the detector. The red distribution shows what the signal of a monoenergetic signal would look like with only energy resolution incorporated into the response function (ignoring effects from quenching and energy loss). The dotted black distribution includes also includes a downward shift to match the response matrix effects, but ignores effects from escaping energy [55].

Chapter 4

PROSPECT IBD Analysis

4.1 Inverse Beta Decay Measurement

After the characterization with all calibration data sets and definition of the detector energy response, the IBD event selection must be optimized to have good background rejection and high statistics of the antineutrino signal. The general concept for identifying IBD events is to find a prompt signal corresponding to the positron emitted by the IBD interaction in coincidence with a delayed neutron capture on ^{6}Li, while rejecting events that may be caused by background events. Using the properties of events and clusters defined in the previous section, several event cuts were individually optimized using studies with independent scans through selection parameters. All studies were done using 20% of the full dataset, with events selected uniformly across the entire time period of the data. In order to optimize the cuts, a number of different figures of merit (FOM) were used including an effective statistics metric and both a global signal-to-background ratio and targeted signal-to-background ratios. An effective counts metric (EC) was used rather than actual counts in order to incorporate the effect of background subtraction:

$$EC = \sum_{0.8\text{MeV}}^{7.2\text{MeV}} \frac{1}{\sigma^2_{\text{relative}}} \tag{4.1}$$

where 0.8-7.2 MeV represents the pre-defined energy range for the IBD analysis in 0.2 MeV wide energy bins and $\sigma_{relative}$ is the relative uncertainty on an energy bin after a time-weighted background subtraction. An additional 5% uncertainty is added to the background uncertainty as a conservative estimate of any issues in the background subtraction.

In addition to this main FOM, signal-to-background ratios were also used as secondary handles on the event selection studies. These metrics are the total integrated rate of the IBD spectrum divided by the integrated rate of the background spectrum. The three different signal-to-background FOMs were integrated over different energy ranges either on the entire energy range (0.8-7.2 MeV), the nH background energy range (1.8-2.2 MeV), or the nC^* background energy range (4.0-5.0 MeV). These last two ranges correspond to the two dominant backgrounds relating to events of a neutron capturing on H and a neutron inelastically scattering on C respectively. While the total signal-to-background FOM helped optimize cuts that offered general improvement across the entire spectrum, the two background specific FOMs offered a more sensitive metric for selection cuts that helped improve just one of these backgrounds.

The following subsections explain the specific cuts and their motivations, including cuts on the prompt event, cuts on the delayed event, cuts on the correlations between the two events, as well as background rejection vetos. After the description of the selection cuts, the analysis is described to produce the IBD spectrum from the dataset.

4.1.1 Event Selection

For the prompt event, the IBD is expected to deposit less than 8 MeV of energy, including the two annihilation gammas produced, within a few nearby segments (multiplicity usually less than 3). The majority of the energy is usually deposited in the segment where the IBD interaction occurred. Because these interactions occur on the nanosecond timescale, the entirety of this will occur in a single, prompt cluster in the detector. Therefore, the

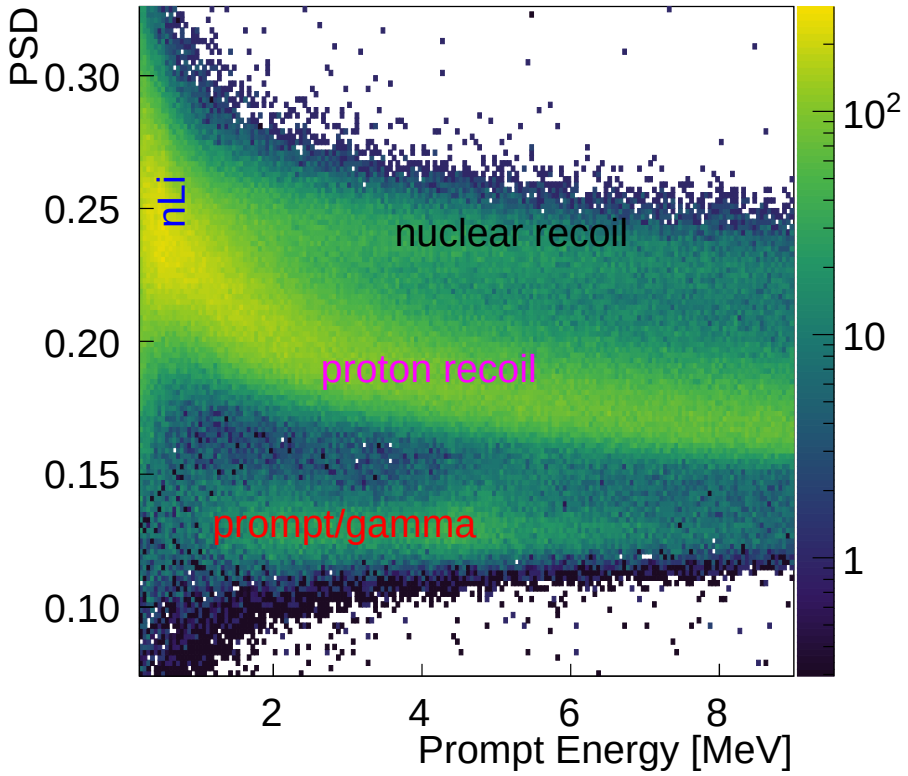

Figure 4.1: 2D histogram of event clusters from one reactor-on cycle at HIFR. The x-axis is reconstructed energy, and the y-axis is the PSD metric. Distinct bands are visible by eye for the prompt/gamma-like events, proton recoils, and nuclear recoils. A peak around 0.5 MeV and 0.25 PSD shows the nLi capture events [55].

initial prompt event of the IBD interaction is selected as an event that with a reconstructed energy between 0.8 and 7.2 MeV. Additionally, a cut on the prompt event PSD is made. For different slices of energy, the PSD distribution is fit with two Gaussians around the gamma and neutron bands. The events all must be within 1σ of the fitted mean of the gamma band to be selected as a prompt event of the IBD. A 2D histogram of all events from one reactor-on cycle is given in Figure 4.1 [55], where the reconstructed energy is given on the x-axis and PSD is given on the y-axis. Besides the loose reconstructed energy cut, and the energy-dependent PSD cut, there are no other cuts specifically on the prompt cluster of the IBD event.

The neutron produced in the IBD thermalizes in the LiLS volume on the 10-100μs timeframe. At that point, approximately 75% of the thermal neutrons capture on a nucleus of ^{6}Li, which then decays into an $\alpha - t$ pair with 0.526 MeV of visible energy. These two products deposit all of their energy within 10's of microns, which is almost always within a single segment. Based on 2D gaussian distribution fitted on the nLi peak in both energy and PSD, the event must be within 3.0 σ in energy from the central fitted value and within 1.8 σ from the fitted PSD mean for the event distribution.

In addition to the energy and PSD selection, the volume is fiducialized to reject events whose primary segment is situated near the edge of the detector, where backgrounds are more prevalent. Any event with a prompt or delayed event primarily in the left-most column, the right-most column, the top row, or the bottom row of the detector is excluded from the analysis. Additionally, due to a hot spot of gamma backgrounds near one part of the detector, two additional segments are excluded from the fiducial volume (coordinates (11,1) and (12,1) using the indices from Figure 3.5).

In addition to selection cuts on single events, the correlations between the prompt and delayed event are also optimized using the same selection criteria as the singles cuts. In the IBD event, the neutron is expected to be captured after the positron event with a mean lifetime of approximately 50 μs. To accommodate a few mean lifetimes, the time coincidence is set so that the neutron-like event must occur between 1 and 120 μs after the electron-like event. The first μs is omitted from the range to alleviate any issues from an overlap between the two events.

In addition to defining this spectrum via this timing correlation, an additional timing window is measured to account for accidental background events. This additional window is measured where nLi events occur 2-12 ms before the electron-like event in time. The spectrum measured in this time window is expected to contain completely uncorrelated events, and is scaled by time in order to directly subtract off the accidental background spectrum from the total measured spectrum in the correlated time window to produce only

the correlated background spectrum for the analysis.

Finally, because the mean path length of the positron and neutron capture products from the primary IBD interaction are order 10cm and 10μm respectively, a spatial correlation cut is placed on the prompt and delayed event. If the events are in the same primary segment, they must fall within 18cm of each other. If the events are in adjacent primary segments, their z_{rec} positions must fall within 14cm of each other. Any other greater distances, or events in non-adjacent segments are cut from the analysis.

The last selection criteria for for the IBD events is that they are clear from any events that are correlated with backgrounds, which is imposed by implementing a number of veto cuts. The first veto implemented is a pileup veto, meaning that all events must have no events within the first 800ns after the prompt event. This cut is made to clean up any pileup events that may lead to complicated overlapping waveforms.

Another veto is the muon veto, stating that within 140μs after the prompt event there cannot be a high energy muon event, defined by any event with a total E_{rec} greater than 15 MeV across all segments. These muon events have a strong correlation with the nH capture background event. A neutron veto is also implemented, where a second nLi capture cannot occur within 400μs before or after the prompt event. This is implemented to remove any degeneracy issues since a single prompt IBD event can only have one nLi capture. The last veto cut is placed such that no events in the neutron band in Figure 4.1 can occur within 250μs after a prompt event. These nuclear recoil events are correlated with the nC* background, and inclusion of this veto increases the signal-to-background ratio around 4.0-5.0 MeV.

After defining these event selection criteria, these cuts are applied to the full data set (both reactor-on and reactor-off periods). The accidental-subtracted event rates (blue) are plotted as a function of time in Figure 4.2 [55]. In this plot, data points correspond to roughly one live-day of data taking show the significant difference in event rate due to reactor operation. This showcases PROSPECT's unprecedented ability to measure an-

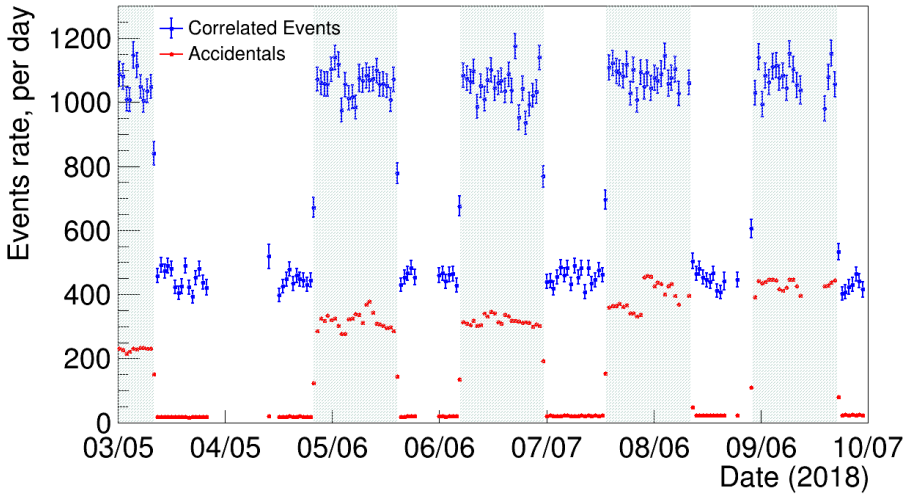

Figure 4.2: Correlated (accidental) event rates over time are in blue (red). Each data point corresponds to one live day of data taking, and showcases the difference in event rates between reactor-on and reactor-off data taking periods, with a signal-to-background ratio greater than 1. Accidental background rates are already subtracted out of the correlated rates [55].

tineutrinos at the Earth's surface with a signal-to-background ratio greater than 1.

4.1.2 Correlated Backgrounds

Once the event selection has been defined, the next goal is to directly subtract out the correlated backgrounds in a data-driven manner. This is primarily done by subtracting out the measured spectrum during reactor-off periods weighted by livetime (which only contain correlated backgrounds) from the measured spectrum during reactor-on periods (which contain both the IBD signal and the correlated backgrounds). The original, total spectrum for the reactor-on data set and reactor-off data set are given in Figure 4.3 [55] for reference. Here the solid blue (red) line gives the total, reactor-on (-off) correlated distribution after accidental subtraction, while the dashed line shows the accidental spectrum that has been subtracted out of the correlated distribution. The accidental spectra have been scaled by the ratio of timing acquisition windows to the correlated distributions, and the reactor-off spectra have also been scaled by the ratio of total livetime to the reactor-on

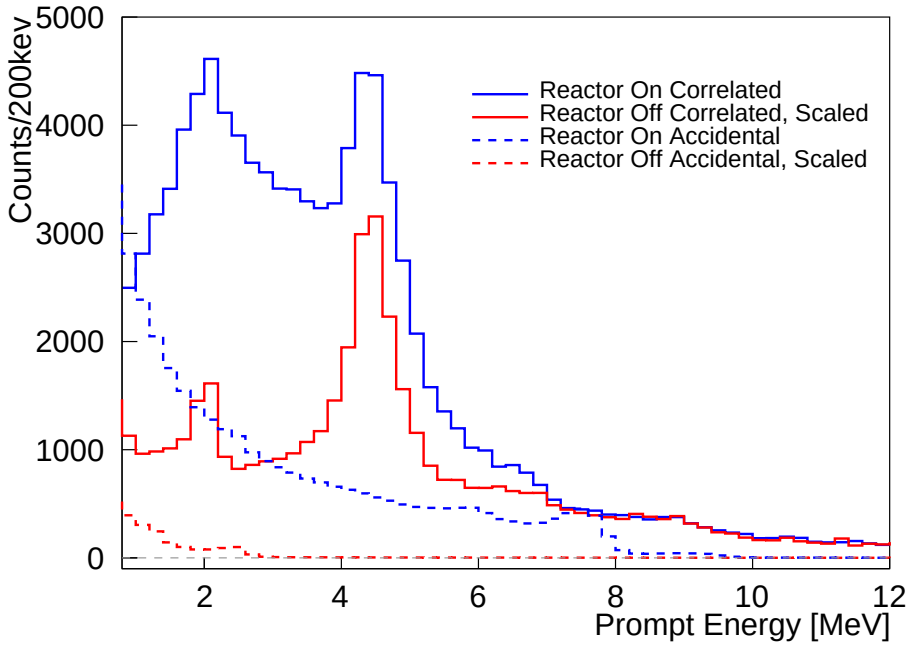

Figure 4.3: Measured E_{rec} for the correlated IBD candidates for both reactor-on (blue) and reactor-off (red)data taking periods. The timing acquisition window scaled accidental backgrounds, which have already been subtracted out, are shown in dashed lines for reference. The accidental spectrum for the reactor-on period is much greater due gamma activity associated with the operation of the HFIR reactor. Previous surveys show no significant change in neutron backgrounds correlated with the reactor [55].

spectra. While the accidental rate is significantly higher during the reactor-on data period, this is associated with gamma singles correlated with reactor operation, but previous background surveys [68] show no significant change in correlated neutron backgrounds in the area from the reactor.

In order to account for variations in cosmogenic backgrounds between reactor-on and reactor-off data taking periods, the correlation between local atmospheric pressure and cosmogenic flux can be used. By directly measuring this correlation between different reactor-off periods, a correction factor for the difference between the reactor-on and reactor-off datasets can be made. The total correlation between the atmospheric pressure and the background rate is shown by the slope of the fitted line in Figure 4.4 [55], with a fitted correlation of -0.036$\pm$0.005 in units of the event rate per pressure [$s^{-1}Bar^{-1}$]. Par-

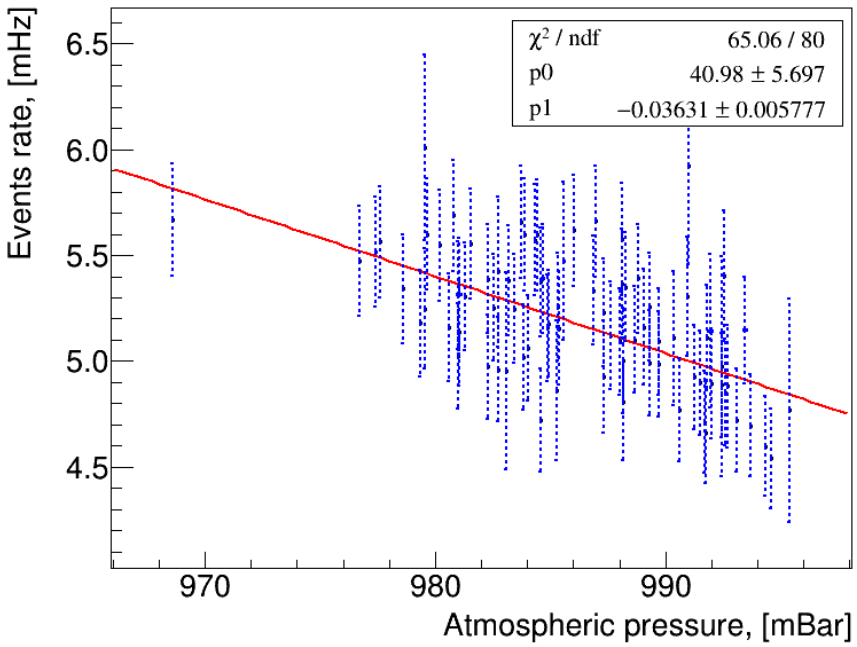

Figure 4.4: Event rate of cosmogenic background events during reactor-off data taking as a function of atmospheric pressure. Each point represents one calendar day of data. The fitted slope is used to produce an aggregate correction factor for correlated background subtraction between reactor-on and reactor-off data taking periods [55].

tially due to the staggering of reactor-on and reactor-off data taking periods, any long term seasonal variations are mitigated, leaving only a small corrective factor of 1.00±0.03% is derived from the difference in average pressure between reactor-on and -off periods.

Besides cosmogenic correlated backgrounds, other possible sources of backgrounds were investigated. The three investigated background sources were neutrinos from spent nuclear fuel, time-coincident backgrounds from gammas and neutrons produced in the reactor, and time-correlated signals from α sources in the PROSPECT detector. Based on calculations of expected interaction rates using nuclear cross sections for materials in the detector and facility, incorporating the geometry of the sources and water pool, these sources are all expected to have event rates $<$0.1 events per day, making a negligible contribution to the final measurement [55]. These factors are not directly incorporated into the analysis

Actual sources of neutrino interactions besides from the direct fission of ^{235}U are accounted for, however. Rather than calculating these rates and subtracting them from the data, however, these contributions are included in the reference model for the spectral analysis, and are described in Section 4.3.

4.1.3 IBD Event Rates

While the general strategy has been outlined in the previous sections, a more thorough description is included here for the formal calculation of the final IBD rates based on the muon vetos. Let $R[x]$ be the event rate of some event x, which in general can be a selection of an event, or a compound event selection with multiple criteria. Additionally let γ, n, and μ represent the distinct event selections for a prompt electron/gamma like event, a delayed nLi capture event, and a muon event respectively as described in the previous subsections. Further, let the symbols $\odot, \oplus$, and $+$ represent two event classes that are truly in correlated coincidence with each other, truly in accidental coincidence with each other, or in coincidence with each other respectively. For example, $R[\gamma + n]$ is the rate of events with a prompt γ event followed by a delayed n event following the selection criteria previously described in the Section 4.1.1.

The incorporation of the muon veto means that ideally we want to find $R[(\gamma \odot n) \odot \not{\mu}]$, which is the event rate of γ events truly correlated with n events that are not truly correlated with μ events. Of course we cannot differentiate in the data for a given event selection $a+b$ what is a truly correlated coincidence from what is a truly accidental coincidence, but we can state that the rates behave as

$$R[a \odot b] = R[a + b] - R[a \oplus b], \tag{4.2}$$

and we can also assume that the truly accidental rate is equal in any selection time window up to a scaling factor based on the ratio of the time intervals. Let us define an operation

for using an offset timing interval to calculate accidental coincidences as $\otimes$, such that $R[\gamma \otimes n] = R[\gamma \oplus n]$ up to a scaling parameter defined by the ratio of the acquisition times.

Then let us also note that for some event rate (for example n) not correlated with a muon, the rate

$$\begin{aligned} R[n \odot \mu] &= R[n + \mu] + R[(n \odot \mu) \oplus \mu^*] \\ &= R[n + \mu] + f_\mu R[n \odot \mu] \end{aligned} \tag{4.3}$$

where μ^* represents at least one muon event in the coincidence timing window and f_μ is a scaling factor that can be directly calculated as the probability of μ^* events falling within the accidental cut based on the measured muon singles event rate. Rearranging Eq. 4.3 yields

$$R[n \odot \mu] = \frac{R[n + \mu]}{1 - f_\mu} \tag{4.4}$$

where this relation extends to any event selection (not just n) being not correlated to a μ event, as long as the timing window is consistently defined. Using Eq. 4.4, and then expanding with Eq. 4.2, our goal IBD rate can be rewritten as

$$R[(\gamma \odot n) \odot \mu] = \frac{R[(\gamma \odot n) + \mu]}{1 - f_\mu} = \frac{R[(\gamma + n) + \mu] - R[(\gamma \oplus n) + \mu]}{1 - f_\mu}. \tag{4.5}$$

Finally, while the first term in the numerator is easily selected on from the data, the second term is not due to the "truly accidental" operator. This can be expanded using the off-time window accidental operator, but must include an additional factor of $(1 - f_\mu)$ to account for a double-counts of accidental vetos appearing in both timing windows, or

$$R[(\gamma \oplus n) + \mu] = \frac{R[(\gamma + \mu) \otimes (n + \mu)]}{1 - f_\mu}. \tag{4.6}$$

Category	Reactor-On	Reactor-Off
Calendar Days	95.65	73.09
Live Days	82.25	65.16
IBD Candidates	115852	30568
Accidental Backgrounds	28358±18	1309±4
Correlated Candidates	87494 ± 341	29258±175
Rate Per Calendar Day	915± 4	400±2
Cosmogenic Backgrounds	36934±221	N/A
Total IBD Signal	50560±406	N/A
Rate Per Calendar Day	529 ±4	N/A

Table 4.1: Statistics of selected IBD candidates and accidental/cosmogenic backgrounds. Errors, where included, represent statistical uncertainties in the relevant signal and background datasets [55].

This combines with Eq. 4.5 yields

$$R[(\gamma \odot n) \odot \not{\mu}] = \frac{R[(\gamma + n) + \not{\mu}]}{1 - f_\mu} - \frac{R[(\gamma + \not{\mu}) \otimes (n + \not{\mu})]}{(1 - f_\mu)^2}. \tag{4.7}$$

From this equation, the total IBD event rate with all cuts, efficiencies, and vetoes accounted for during all reactor-on and reactor-off data taking periods. The total event counts and live time rates are included in Table 4.1 [55], with a total IBD signal of 50,560 ± 406 events, with a total rate per calendar day of 529 ±4 events. Additionally, the IBD rates for each segment (including an efficiency correction from MC studies) is shown in Figure 4.5 [55]. The relatively large size of the PROSPECT detector relative to the baseline from the reactor shows the dependence of the neutrino flux on the inverse of the baseline squared. In this plot, the IBD rate is integrated over the main analysis region of 0.8 to 7.2 MeV. An inverse square function is fit to the data (the only parameter being a total normalization constant), yielding a χ^2/dof of 72.4/69 showing excellent agreement with the expected baseline dependence.

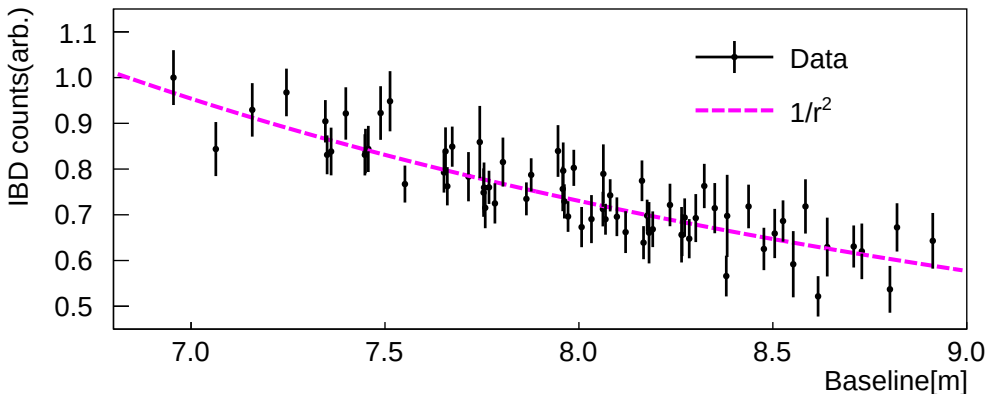

Figure 4.5: Efficiency corrected IBD rate integrated from 0.8 to 7.2 MeV for each segment as a function of baseline from the reactor core. An inverse square function is fitted (normalization parameter only) to the data to show the expected baseline dependence of the rates. The fit yields a χ^2/dof of 72.4 / 69, showing excellent agreement between the data and the expected baseline dependency [55].

4.2 Oscillation Analysis

4.2.1 Data Sets and Predictions

For the sterile neutrino search, the relative spectral differences between different baselines are used to search for signatures of sterile neutrino oscillations. In order to obtain roughly equal IBD rates in each baseline bin, the baseline binning is defined as shown in Figure 4.6 [55]. Note that due to the inverse square dependence of the antineutrino flux, there are fewer segments in the near baseline bins than in far baseline bins (baseline bins are indexed in order of increasing baselines in the figure). Across the full data set, there are roughly 5000 events per baseline bin, with per-bin relative variation of order 10%.

The full data set is binned both by baseline l and energy e, ranging from 6.7 - 9.3 m and 0.8 - 7.2 MeV respectively, with bins denoted as $M_{l,e}$. The full data is presented in Figure 4.7[55], which includes the measured spectra at each baseline in black as well as the PG4 generated prediction assuming no sterile neutrino mixing in cyan. The prediction values $P_{l,e}$ are based off of the Huber model [44], which is passed through a detector response

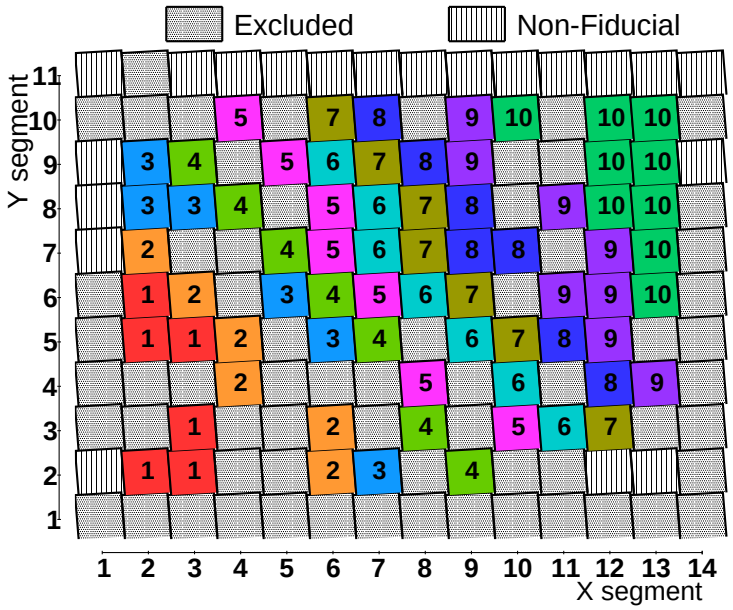

Figure 4.6: Baseline binning of segments for the oscillation analysis. Bins are defined to have roughly equal total IBD counts per bin, where the inverse square dependence of baseline of the IBD flux means there are fewer bins at near baselines than far baselines. Only fiducial, active active are sorted into baseline bins [55].

function for each baseline bin. These response functions incorporate segment-to-segment variations due to inefficiencies from each segment's position relative to inactive material, such as the edge of the detector or inactive cells. Predicted baselines also incorporate both the non-zero volume of the reactor core as well as the limited positional reconstruction in the dimensions of the optical array (x and y dimensions). A common normalization factor is included in the prediction generation set to match the total measured IBD signal count.

The prediction values $P_{l,e}$ can all be distorted by a possible oscillation to a sterile neutrino flavor before application of the detector response function. In addition to the particular baselines l and energies e, this distortion is parameterized by the sterile neutrino mixing parameters $\Delta m^2_{41}, \sin^2\theta_{14}$ (similarly defined as the parameters in Section 1.3). Ultimately the comparison of data to these predictions is used to search for sterile neutrino signatures. To remove any first order dependence of these predictions on the initial Huber model, each baseline spectrum is compared to a prediction scaled by the total spectral

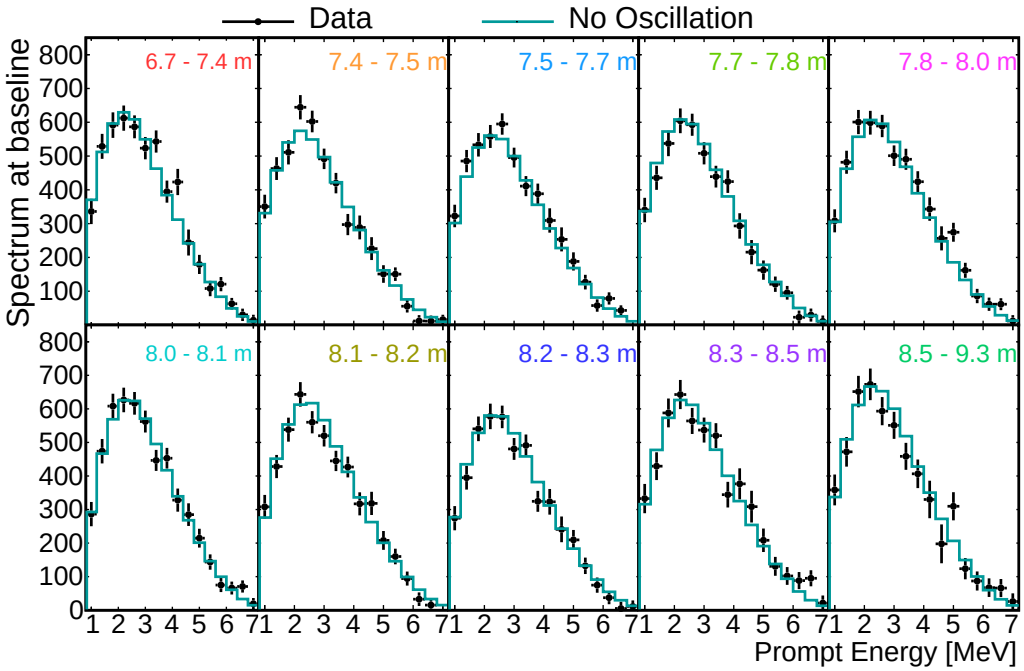

Figure 4.7: PROSPECT data E_{rec} spectrum across 10 baselines used for the oscillation analysis. Data is given in black, with the predicted spectrum at each baseline assuming no sterile neutrino mixing. All error bars plotted are statistics-only. For the prediction a common global normalization factor is included to match total IBD signal counts, and relative normalizations are included based off of segment-to-segment relative efficiencies for each baseline as calculated from simulation [55].

measurement. That is, each measurement $M_{l,e}$ it is compared to $M_e \frac{P_{l,e}}{P_e}$ where

$$M_e = \sum_{l=1}^{10} M_{l,e} \text{ and } P_e = \sum_{l=1}^{10} P_{l,e} \tag{4.8}$$

where the ratio $\frac{P_{l,e}}{P_e}$ reduces the model dependency. Perhaps more easily illustrated in Figure 4.8 [55], the ratio of $M_{l,e}$ to $M_e \frac{P_{l,e}}{P_e}$ will be consistent with unity for all energies and baselines in the absence of sterile neutrinos. For reference, the prediction using the RAA best fit sterile neutrino parameters [34] is included for reference, as is the prediction corresponding to the best fit from the following analysis.

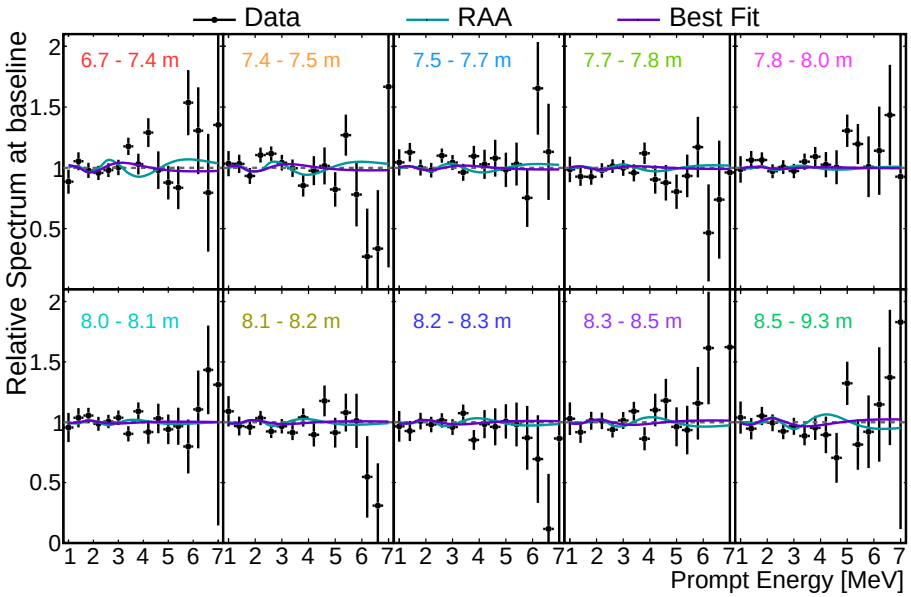

Figure 4.8: Ratio of data to prediction as scaled by the ratio of the integrated prediction to the integrated data ($\frac{M_{l,e}}{M_e} \times \frac{P_e}{P_{l,e}}$) at each of the 10 PROSPECT baselines. The dashed line at unity represents the prediction in ratio space for the case of no sterile neutrino, while the cyan line corresponds to the prediction given the RAA best fit, with the purple line showing the result from the best fit. All error bars plotted are statistics only [55].

4.2.2 Uncertainties

For the oscillation analysis, in order to maintain better statistics per bin the energy is binned in 0.4 MeV wide bins (rather than the 0.2 MeV wide bins in the spectral analysis), which yields 16 energy bins. This, alongside the 10 baseline bins, make Δ a 160 element vector of the differences between the measurements and relatively scaled predictions for each energy and baseline bin e and l:

$$\Delta_{l,e} = M_{l,e} - M_e \frac{P_{l,e}}{P_e}. \tag{4.9}$$

These elements are indexed in increasing energy bins within the closest baseline before restarting at the lowest energy bin in the next closest baseline bin, etc. In order to quantitatively test the sterile neutrino hypothesis, a χ^2 test statistic is defined using a covariance

matrix approach

$$\chi^2 = \Delta^{\mathrm{T}} V_{tot}^{-1} \Delta, \tag{4.10}$$

where V_{tot} is a 160×160 element total covariance matrix for the measurement. V_{tot} is composed of the sum of the statistical and systematic covariance matrices V_{stat} and V_{sys}. V_{stat} is a mostly diagonal matrix dominated by the reactor-on IBD candidates, including effects from correlated background subtraction for each bin handled using a Poisson-based error propagation, including scaling for relative livetime corrections between data taking periods. Accidental background subtraction is also included in this error calculation, but it has very little contribution to the total uncertainty due to the large offset timing window used to acquire the accidental spectrum. These uncertainties are largely uncorrelated between energy and baseline bins, but do include some statistical correlations on the off diagonals since M_e is dependent on each $M_{l,e}$.

V_{sys} is a more complicated covariance matrix than V_{stat} and is calculated through dedicated studies of several components, each calculated as a relative uncertainty. In these studies, a nominal uncertainty is conservatively defined for each parameter. Then 1000 toys are generated for the oscillation relative spectra sampling from a Gaussian distribution about that parameter based on the nominal uncertainty where the parameter may or may not be correlated across energies and/or baselines depending on the systematic uncertainty being studied. The full 160×160 covariance matrix is then calculated comparing each toy vs the model spectra as generated without any uncertainty on the given parameter. In order to remove any model dependence, these covariance matrices are then divided out by the model spectrum in order to produce a relative covariance matrix for the parameter. The separate systematic covariance matrices are added together and then scaled by the final measurement to have the appropriate statistical magnitude for the measured data. That is, each element $V_{sys,rel}^{i,j}$ is multiplied by $M^i \times M^j$ to scale to the total systematic uncertainty in order to minimize any model dependence used in each study (though all systematic

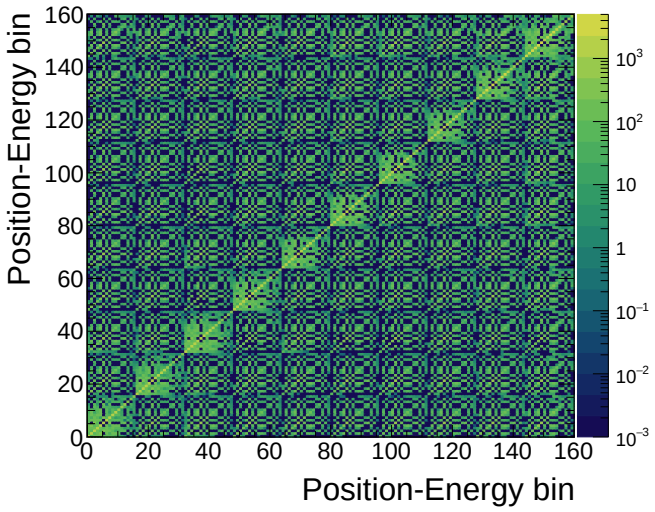

Figure 4.9: Total covariance matrix used for the oscillation analysis, containing all correlated and uncorrelated statistical and systematic uncertainties. Statistics are dominated by uncorrelated uncertainties of the IBD spectrum including correlated background subtraction, and systematic uncertainties are generated by dedicated MC toy studies modeling uncertainties with Gaussian errors [55].

studies use the Huber model as an underlying model). The specific uncertainties studied relate to effects from the detector response, detector stability in time and across segments, and background estimates, and the end result, V_{tot} is given in Figure 4.9 [55]. Details of the uncertainties are summarized in Table 4.2 [55], with more detailed descriptions for each as follows:

- *Absolute background normalization*- due to possible issues in the absolute normalization of the correlated background spectrum, whether due to improper atmospheric scaling or unforeseen issues related to the difference between reactor-on and reactor-off data taking periods, a relative uncertainty dependent on the normalization of the total reactor-off correlated spectrum is accounted for correlated across all energies and baselines.

- *Absolute nH peak normalization*- Similar to the total background, an additional normalization uncertainty is ascribed to the nH peak in the background due to possibly

different dependencies on that background from other reactor-off backgrounds. This uncertainty is correlated across all baselines, but only across energy bins associated with the nH background range from 1.8-2.2 MeV. This uncertainty is treated as uncorrelated with the total background normalization.

- *Relative signal normalization*- In order to account for possible variations between the size and efficiency between optical segments in different baselines, a relative normalization uncertainty is incorporated which is correlated across energies but is independent across baseline bins.

- *Baseline*- To account for possible errors in the position of the detector relative to the core of the reactor, which impacts the calculation of the absolute baseline values, a parameter that is correlated across all energies and all baselines is varied to evaluate it's impact on the analysis.

- *Energy scale non-linearity constants*- The three parameters of the energy scale non-linearity are independently varied based on their uncertainties. In particular, the first- and second-order Birks constants and the Cherenkov contribution to the energy scale uncertainties contribute to the uncertainties and are correlated across all baselines.

- *Energy scale linearity*- To account for total linearity differences in the energy scale of the detector, whether due to changes across the entire detector in time or on a segment-by-segment basis, uncertainties both correlated across baseline as well as uncorrelated across baselines is included based off of stability cross checks in time and across segments.

- *Photostatistics resolution*- Possible differences either across time or segments between the collection of optical photons could contribute to systematic uncertainties. These could be identical across the different baselines of the detector (such as in changes in the optical properties of the LiLS) or uncorrelated between baselines

(such as a damaged reflector panel or other part of the optical path in a particular segment), and thus there are covariance matrices generated both correlated across baselines and uncorrelated.

- *Energy loss and leakage*- Energy loss and leakage can occur either as a correlated effect across baselines or as an uncorrelated effect. Improper modeling of the energy deposited in inactive segments could result in either baseline correlated or uncorrelated effects depending on the modeling issue while variances in the thickness of reflector panels, which are dead volume in the detector, could result in uncorrelated effects as well. All three of these mechanisms are used to generate separate covariance matrices.

- *Energy analysis threshold*- While there is a hardware-level threshold set in the DAQ trigger subject to the relative gain differences, a higher threshold is set in the analysis so that thresholds can be set by reconstructed energy rather than uncalibrated raw signals. Regardless, differences in energy calibration, either in time or across segments, could lead to differences in total E_{rec}. These mechanisms are treated both as baseline correlated and uncorrelated uncertainties respectively.

4.2.3 Test Statistics and Results

With the data, predictions, and uncertainties defined in the previous subsections, as well as the χ^2 defined in Equation 4.10, the search for a sterile neutrino signature can begin. By calculating the χ^2 value from predictions $P_{l,e}$ across an array of different sterile oscillation parameters Δm^2_{41} and $\sin^2\theta_{14}$, the difference between χ^2 from the best fit value can be used to produce a $\Delta\chi^2$ map, which is given in Figure 4.10 [55].

For the given χ^2 test statistic, it is clear to see the preference for a non-zero sterile neutrino, but confidence of that preference over a null oscillation or the RAA best fit point is

Parameter	Nominal Value	Uncertainty	Correlations
Absolute background normalization	-	1.0%	Correlated between energies and baselines
Absolute n-H peak normalization	-	3.0%	Correlated between energies and baselines
Relative signal normalization	-	5%	Correlated between energies
Baseline uncertainty	-	10 cm	Correlated between energies and baselines
First-order Birks constant	0.132 MeV/cm	0.004 MeV/cm	Correlated between baselines
Second-order Birks constant	0.023 MeV/cm	0.004 MeV/cm	Correlated between baselines
Cherenkov contribution	37%	2%	Correlated between baselines
Absolute energy scale	-	0.6%	Correlated between baselines
Absolute photostatistics resolution	-	5%	Correlated between baselines
Absolute energy leakage	-	8 keV	Correlated between baselines
Absolute energy threshold		5 keV	Correlated between baselines
Relative energy scale	-	0.6%	Uncorrelated between baselines
Relative photostatistics resolution	-	5%	Uncorrelated between baselines
Relative energy leakage	-	8 keV	Uncorrelated between baselines
Relative energy threshold	-	5 keV	Uncorrelated between baselines
Reflector panel thickness	1.18 mm	0.03 mm	Uncorrelated between baselines

Table 4.2: Summary of systematic uncertainties incorporated into the oscillation systematic covariance matrix V_{sys}. Nominal parameter values are provided, where applicable, as well as relevant correlations [55].

difficult to quantify using only the $\Delta\chi^2$ map. In many cases, a standard $\Delta\chi^2$ as compared to the degrees of freedom approach would be sufficient, but the possibility of the best fit being on the boundary of the allowed parameter space at $\sin^2\theta_{14} = 0$ violates the assumptions of Wilks' theorem, an underlying theorem used to approximate the $\Delta\chi^2$ distribution as an analytic χ^2 distribution to evaluate a confidence level for the measurement. Instead, two independent methods of quantifying the confidence of a sterile neutrino hypothesis are employed: the Feldman-Cousins method [77] and the Gaussian CL_s method [78].

For Feldman-Cousins, a frequentist statistical approach, distributions of the $\Delta\chi^2$ test statistic are generated with 1000 toys per sterile neutrino hypothesis and compared to the null-oscillation. To generate each toy, a 160-element vector of Gaussian-distributed random numbers ($\mu = 0, \sigma = 1$) are multiplied by a Cholesky decomposition of the total covariance matrix, asserting that all toys have correlated and uncorrelated uncertainties consistent with the PROSPECT experiment. This random vector is then added to a prediction of the spectrum created by multiplying the predicted spectrum at each baseline by the full detector response matrix. By doing this for 1000 toys, and comparing each of those to the prediction from the null-oscillation hypothesis, a distribution of $\Delta\chi^2$ values can be de-

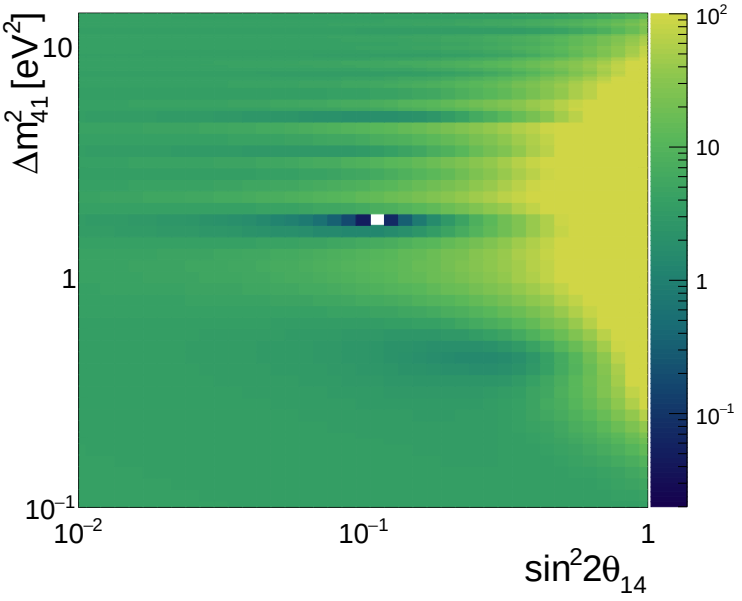

Figure 4.10: $\Delta\chi^2$ map showing the difference between the χ^2 test statistic defined in Equation 4.10 of a given sterile neutrino hypothesis with Δm^2_{41} and $\sin^2\theta_{14}$ and the value from the best fit point at $\Delta m^2_{41} = 1.78\text{eV}^2$ and $\sin^2\theta_{14} = 0.11$ [55].

termined for each point in the phase space explored for sterile neutrinos. For each of those distributions, the critical test statistic value $\Delta\chi^2_{crit}$ can be determined for a given p-value (nominally 95% confidence in this case) at each sterile neutrino hypothesis. Examples of these $\Delta\chi^2$ distributions are given in Figure 4.11 [55] for both the null hypothesis and at the parameters for the RAA best fit hypothesis, along with red lines denoting the actual $\Delta\chi^2$ from data. For these particular hypotheses, PROSPECT yields p-values of 0.57 and .015 respectively, showing that the null oscillation is not excluded and the RAA hypothesis is excluded at 2.5σ. It should be noted that simply using an analytic χ^2 distribution rather than this Feldman-Cousins generated distribution can lead to notable misinterpretations of a result. Figure 4.12 [79] shows an example of the analytic and numerically generated $\Delta\chi^2$ distributions and their corresponding 1, 2, and 3σ thresholds for the null oscillation. This illustrates the importance of properly modeling these test statistic distributions and why utilizing Wilks' Theorem in this situation is not reasonable.

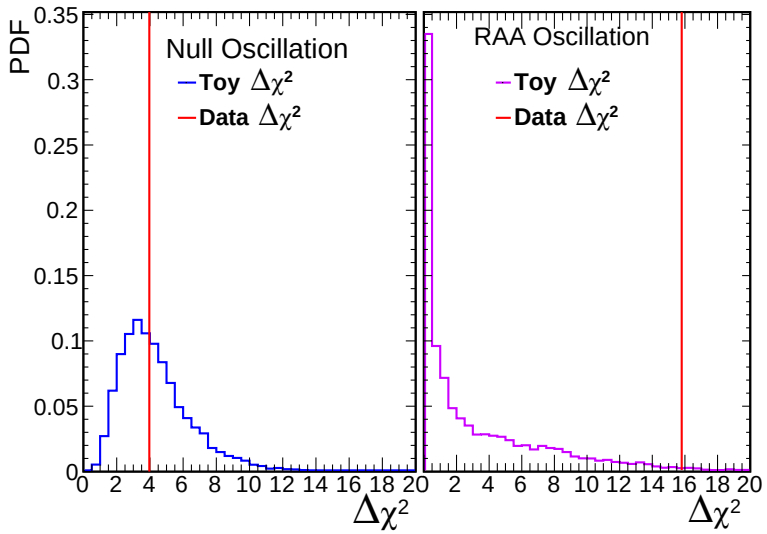

Figure 4.11: Toy MC generated $\Delta\chi^2$ PDF distributions for the null oscillation (left) and RAA best fit (right) hypotheses. Red vertical lines show the $\Delta\chi^2$ value from the PROSPECT data. Integrating to the right of the red line, the p-value for a particular hypothesis can be directly calculated without having to invoke Wilks' Theorem [55].

Rather than just calculating the $\Delta\chi^2$ distribution for two specific hypotheses, this process is applied to all points in a grid scan of the (Δm^2_{41},sin$^2\theta_{14}$) phase space. Then by evaluating the $\Delta\chi^2$ at each point with respect to the $\Delta\chi^2_{crit}$ for that hypothesis, an exclusion curve (or in principle a confidence interval) is drawn. That, in addition to sensitivity curves as calculated from averaging the results from 1000 toy datasets, is given in Figure 4.13 [55]. Also included is the result from the RAA best fit point for reference.

In addition to the Feldman-Cousins, a Gaussian CL$_s$ approach is employed as a cross-check. For the CL$_s$ method, the likelihood of a pair of hypotheses are compared against each other, namely a null oscillation and a hypothesis for a given point in the sterile neutrino mixing parameter space (Δm^2_{41},sin$^2\theta_{14}$). Since this only compares the relative likelihoods of one hypothesis over the other, this technique can not be used to establish confidence intervals for a positive discovery and is only able to set exclusion limits (which is reasonable given the result from the Feldman-Cousins analysis). For this technique, we

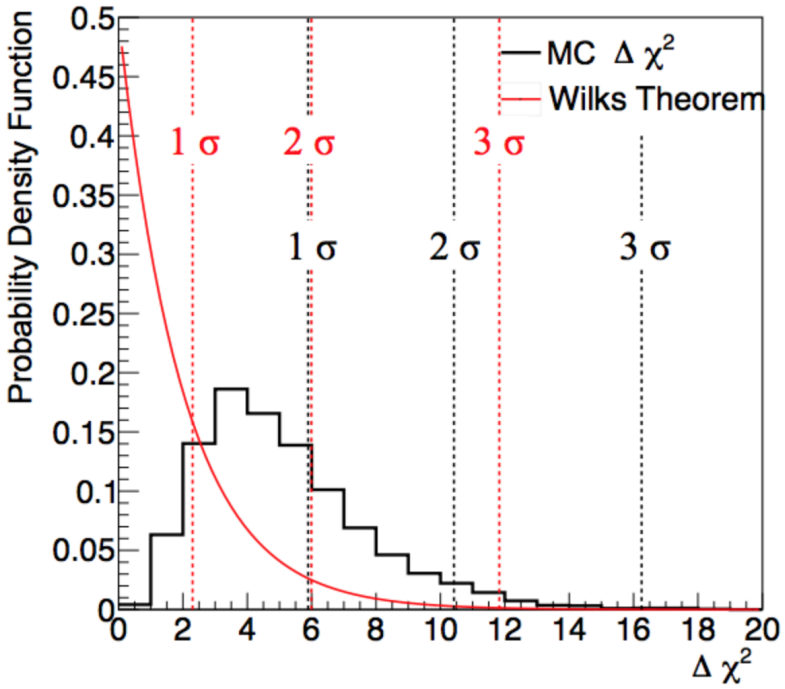

Figure 4.12: $\Delta\chi^2$ distributions calculated analytically via Wilks' Theorem (red) and via MC toys as prescribed by the Feldman-Cousins frequentist approach (black) for the same underlying hypothesis test. The difference between the 1, 2, and 3σ thresholds for the same data highlights the importance of properly modeling the test statistic distribution [79].

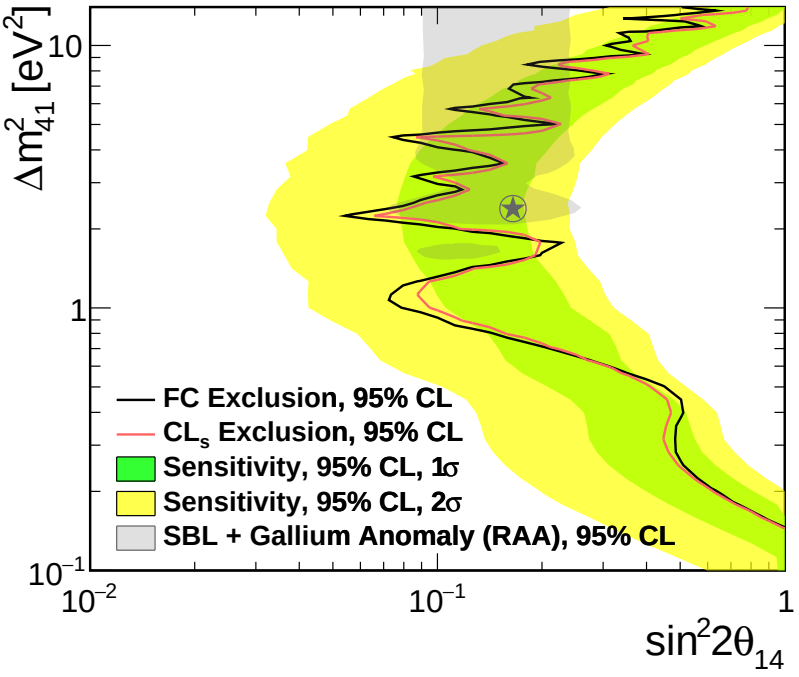

Figure 4.13: PROSPECT sterile neutrino exclusion contour as calculated using the Feldman-Cousins frequentist approach as well as using the Gaussian CL_s technique, the results of which are in agreement. Also featured are the 1 and 2σ 95% CL sensitivities and the RAA best fit point, which is excluded at 2.5σ [55].

define the test statistic ΔT as follows:

$$\Delta T(x) = \chi^2_{min}(x)_1 - \chi^2_{min}(x)_0 \tag{4.11}$$

where $\chi^2_{min}(x)_i$ is the minimum value of the χ^2 defined in Equation 4.10 for hypothesis i, where x represents the measured data, $i = 0$ is the null hypothesis, and $i = 1$ is the alternate hypothesis in question. A positive value for $\Delta T(x)$ represents a preference of the null hypothesis over the alternate hypothesis given data x, and a negative value represents the opposite preference (although the level of preference is not clear at this stage). Note that this definition of ΔT is different from a typical $\Delta\chi^2$ which compares the entire parameter space with respect to the null hypothesis. Because this definition gets the minimum value for each hypothesis independently, it is directly comparing the best fits given a specific pair of hypotheses. Additionally, one can calculate the expected value of ΔT_i using an Asimov data set x_i^{Asimov} where an Asimov data set represents the mean measurement given a hypothesis i. Used for an individual hypothesis (such as in the CL_s method), an Asimov data set can significantly cut down on computational time over approaches like Feldman-Cousins. Specifically, this technique yields:

$$\begin{aligned}
\overline{\Delta T_{H_0}} &= \Delta T(x_{H_0}^{Asimov}) \\
&= \chi^2_{min}(x_{H_0}^{Asimov})_1 - \chi^2_{min}(x_{H_0}^{Asimov})_0 \\
&= \chi^2_{min}(x_{H_0}^{Asimov})_1 \\
\overline{\Delta T_{H_1}} &= \Delta T(x_{H_1}^{Asimov}) \\
&= \chi^2_{min}(x_{H_1}^{Asimov})_1 - \chi^2_{min}(x_{H_1}^{Asimov})_0 \\
&= -\chi^2_{min}(x_{H_1}^{Asimov})_0
\end{aligned}$$

where the terms $\chi^2_{min}(x_{H_i}^{Asimov})_i$ go to zero since the Asimov data set completely aligns with the model, giving a χ^2 of zero by definition. Using the values of the two expected $\overline{\Delta T}$

values as well as the $\Delta T(x)$ from the measurement, one can approximate the distribution of $\Delta T(x)$ as a Gaussian with mean $\overline{\Delta T_{H_i}}$ and standard deviation $2\sqrt{\overline{|\Delta T_{H_i}|}}$ [78] assuming that the true model is H_i.

Given these approximations, one can straightforwardly calculate the exclusion level of a particular hypothesis given the data using the $\mathrm{CL}_s(x)$ value

$$\mathrm{CL}_s(x) = \frac{1-p_1}{1-p_0} \tag{4.12}$$

where $1-p_i$ represents the probability that given hypothesis i, a repeat measurement would yield a greater ΔT value. Specifically, a CL_s close to zero favors H_0 over H_1, but a higher CL_s does not necessarily indicate a preference for H_1. This is why the Gaussian CL_s technique is not useful for setting confidence intervals of positive results, but can only be used to set exclusion limits. The advantage of this method is that its approximation of its test statistic as a Gaussian doesn't invoke Wilks' Theorem and therefore can be accurately applied in this case. Explicitly, the approximation for $1-p_i$ is as follows:

$$1-p_i \approx \frac{1+\mathrm{Erf}\left(\frac{\overline{\Delta T_{H_i}}-\Delta T(x)}{\sqrt{8\sqrt{\overline{|\Delta T_{H_i}|}}}}\right)}{2}, \tag{4.13}$$

and it follows that

$$\mathrm{CL}_s \approx \frac{1+\mathrm{Erf}\left(\frac{\overline{\Delta T_{H_1}}-\Delta T(x)}{\sqrt{8\sqrt{\overline{|\Delta T_{H_1}|}}}}\right)}{1+\mathrm{Erf}\left(\frac{\overline{\Delta T_{H_0}}-\Delta T(x)}{\sqrt{8\sqrt{\overline{|\Delta T_{H_0}|}}}}\right)}. \tag{4.14}$$

Using this metric, any hypothesis with sterile mixing parameters (Δm^2_{41},$\sin^2\theta_{14}$) will be excluded at 95% if the calculation of CL_s for that hypothesis is less than 0.05.

From this definition, the same parameter space is scanned as was done for the Feldman-Cousins analysis, and the Gaussian CL_s exclusion curve is generated. This curve is included in Figure 4.13, and this independent method shows excellent agreement with the

result from the Feldman-Cousins technique. Both of these results yield an exclusion curve that is mostly within the green 1σ sensitivity curve, indicating that these are typical results for the PROSPECT measurement which exclude a large portion of the allowed phase space for sterile neutrino oscillations at 95% confidence.

4.3 Spectrum Analysis

Besides the oscillation search done using a relative spectrum analysis, the total spectrum measured across the entire detector is also investigated. Because more than 99% of the $\bar{\nu}_e$ produced at HFIR are due to ^{235}U, this provides a handle to both compare the data directly against model predictions and evaluate the contribution of this component to the previously measured deviations at LEU reactors [41, 42, 80–82]. With a total of 50560$\pm$406(stat) IBD events and a cosmogenic (accidental) signal-to-background ratio of 1.4 (1.8) this is the highest statistics pure ^{235}U measurement to date.

4.3.1 Non-^{235}U Corrections

For model prediction comparison, the Huber model described in Section 1.5.3 is used as the baseline prediction. In addition to the pure ^{235}U spectrum, an additional correction of less than 1% of the total neutrino flux is independently modeled and added to the Huber spectrum. Contributions from ^{28}Al in the fuel cladding of the core, ^{6}He generated in the beryllium neutron reflector around the core [58], and the average effect from non-equilibrium isotopes produced during the 24-day cycle of HFIR are included in the model prediction. The two materials-based contributions are modeled using the Monte Carlo particle transport code MCNP [59]. The production and $\bar{\nu}_e$ emission of the non-equilibrium contribution is calculated following the procedure in [48]. These contributions are then incorporated into the Huber flux model and convolved with the IBD cross-section described in Section 1.5.3 and then convolved through the detector response function as well. The

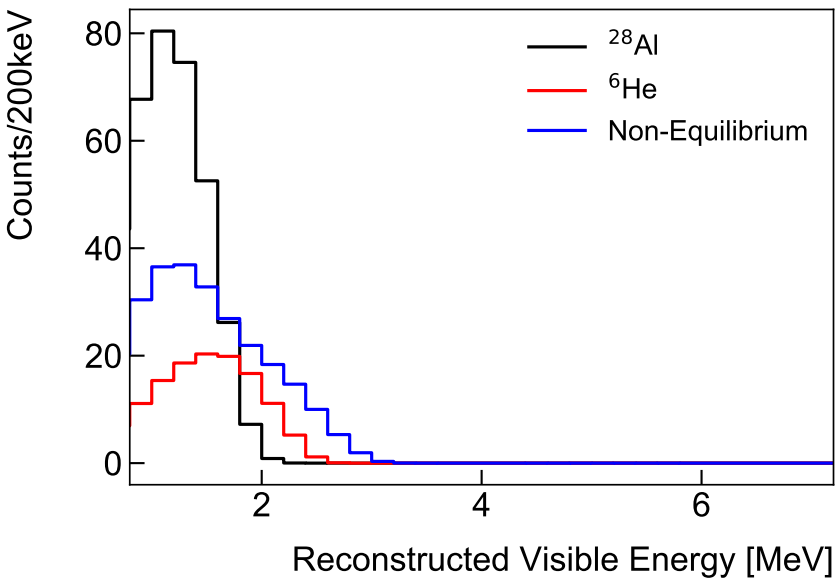

Figure 4.14: Non-^{235}U contributions to the IBD spectrum model. Contributions come from ^{28}Al in the reactor fuel cladding, ^{6}He from the beryllium reflector, and reactor cycle-averaged non-equilibrium isotopes. Total contributions account for $<1\%$ of the total IBD flux and are at relatively low energies [55].

individual contributions are all at low energies (<3 MeV) and are given in Figure 4.14 [55].

4.3.2 Spectrum Uncertainties

The uncertainties relevant to the spectrum analysis are divided into three categoreis. Besides the uncertainties on the Huber model prediction [44], there are detector uncertainties and statistical uncertainties. The detector uncertainties are identified nearly identically to those described in Section 4.2.2 except that they only apply to the total spectrum integrated across the entire detector with no baseline dependence. In particular, most of the uncertainties incorporated into the spectrum analysis correspond 1-to-1 with the baseline-correlated uncertainties given in Table 4.2. The only difference is that an additional 100% uncertainty on the non-^{235}U corrections is incorporated as a conservative value. The statistical detector uncertainties are based on the counting statistics of the total correlated

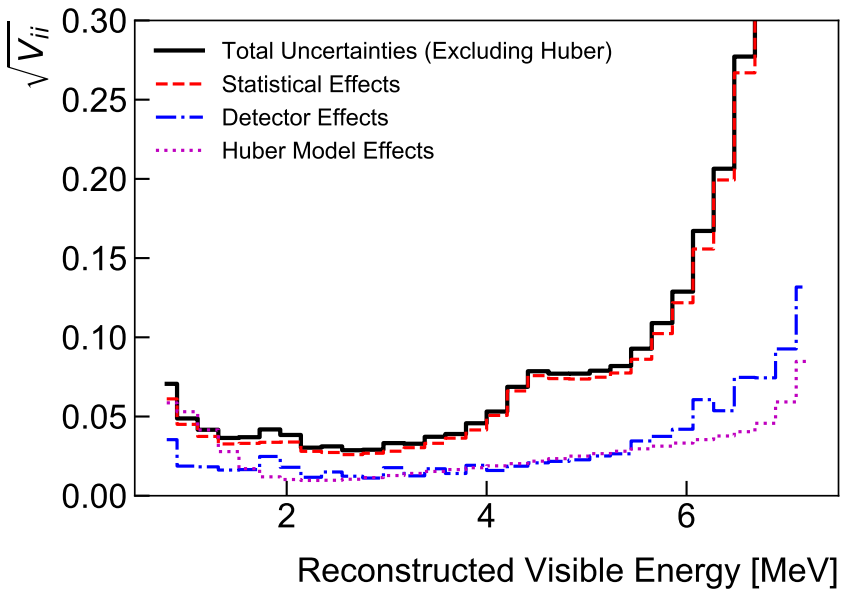

Figure 4.15: Relative uncertainties used in the spectrum analysis, with total model and experimental uncertainties including the breakdown of the statistical and detector effects [55].

background-subtracted spectrum as defined in Section 4.1.3.

After all contributing covariance matrices are generated using the same procedure as before, the matrices are summed to generate the total spectrum covariance matrix V_{tot}. While V_{tot} contains the full contribution from all three categories of uncertainty for the spectral analysis, the contribution from model uncertainties is not included in the uncertainties of the measurement itself in either the error bars in figures or as an input to the joint analysis later in Section 5. The total and individual contributions from the model, statistical, and systematic uncertainties are shown in Figure 4.15 [55], where the relative uncertainty is calculated as the square root of the diagonal of the covariance matrix divided by the measured spectrum. The corresponding breakdown of the components of the systematic uncertainty is included in Figure 4.16 [55].

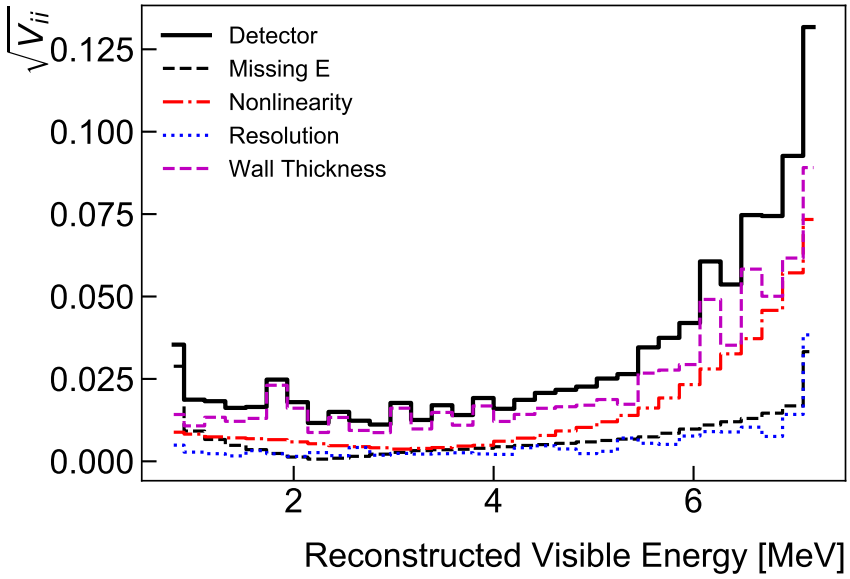

Figure 4.16: Total detector uncertainties and its breakdown into its various components. The total uncertainty of the measurement is statistics dominated [55].

4.3.3 Test Statistics and Results

Given the total covariance matrix V_{tot}, a χ^2 test statistic is calculated to evaluate the agreement between the PROSPECT measurement and the Huber model prediction as follows:

$$\chi^2_{global} = \Delta^T V_{tot}^{-1} \Delta, \tag{4.15}$$

$$\Delta_i = M_i - P_i \times (1 + \eta), \tag{4.16}$$

where M is the measured spectrum, P is the predicted spectrum, and η is a freely floating normalization parameter over which χ^2_{global} is minimized. Here both M and P range from 0.8 - 7.2 MeV in 0.2 MeV bins expressed in counts per MeV. No absolute normalization is included as part of the measurement. This global comparison to Huber yields a χ^2/dof of 30.79/31, corresponding to a p-value of 0.48, showing overall consistency with the model. The full measured spectrum and it's comparison to the model is given in the top panel of Figure 4.17 [55].

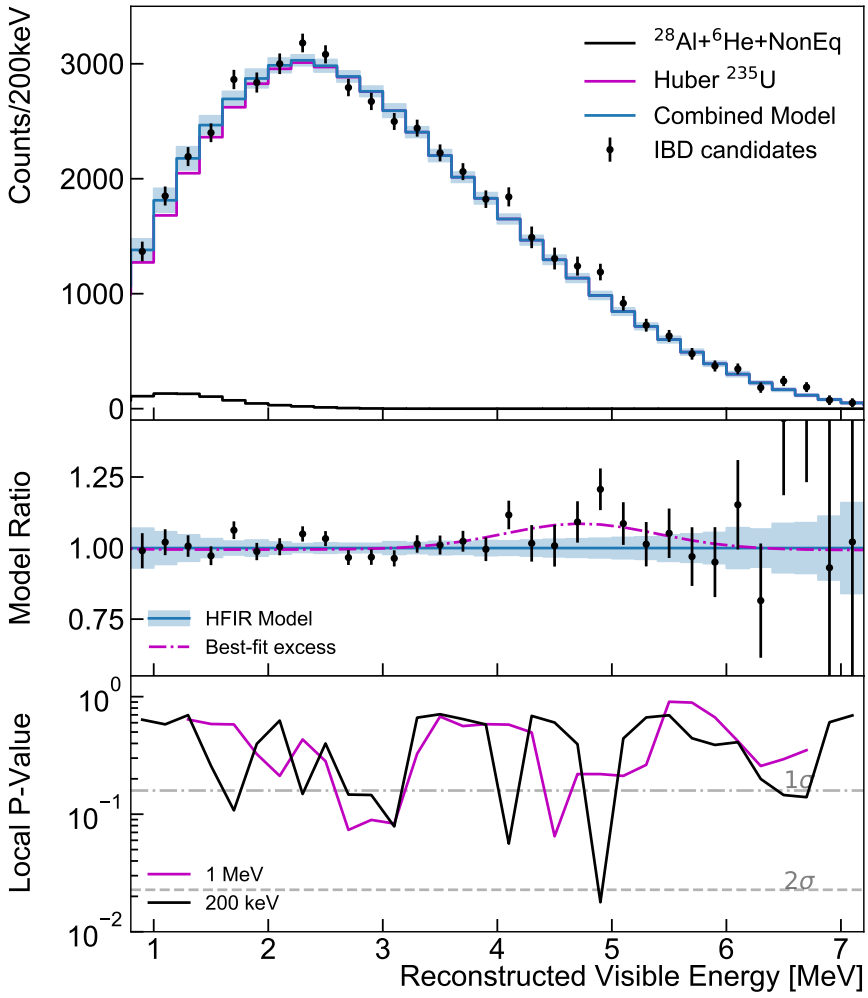

Figure 4.17: (top) ^{235}U spectrum in reconstructed energy as measured by PROSPECT compared to the predicted model. (middle) Ratio of measurement to best fit model, including the best fit excess modeled with a Gaussian with fixed width and mean. (bottom) local p-value in one-bin and 5-bin sliding windows. Error bands on data represent statistical uncertainties, with the band on the model representing model uncertainties [55].

In order to evaluate local deviations between the measurement and the prediction, a sliding window approach is introduced. By adding an extra free parameter on one bin in the spectrum and re-minimizing a new test statistic $\chi^2_{local,1}$, this nested hypothesis produces a new lower χ^2 value such that the difference between the new and old test statistic represents the local contribution to the original χ^2_{global} value. For any specific energy bin i where the free parameter is added, this yields

$$\Delta\chi^2_{local,1,i} = \chi^2_{global} - \chi^2_{local,1,i} \tag{4.17}$$

where a local significance is found by calculating the p-value from the $\Delta\chi^2_{local,1,i}$ given one additional degree of freedom. Calculating this for each bin in the spectrum produces the black local p-value curve in the bottom panel of Figure 4.17 [55]. Most local fluctuations are $\sim 1\sigma$ in scale, with a singular bin reaching $\sim 2\sigma$, which is consistent with statistical fluctuations only.

Instead of adding an extra fit parameter onto exactly one bin in the spectrum, a 1 MeV wide sliding window is made by adding independent fit parameters to five adjacent bins at a time. Then by taking a different test statistic $\Delta\chi^2_{local,5,i} = \chi^2_{global} - \chi^2_{local,5,i}$, where i corresponds to the middle energy bin in a five-bin wide window, a local p-value is calculated using with five degrees of freedom to evaluate contributions to the χ^2_{global} from features that are 1 MeV wide in scale. This curve is plotted in magenta in the bottom panel of Figure 4.17, and while some local deviations correlate between the 0.2 and 1.0 MeV wide windows, there are no features with this sliding window with local deviations greater than 1.5σ.

Lastly, in order to directly compare to the spectrum as measured by Daya Bay, the excess in [40] is modeled as a Gaussian distribution with a fixed mean of 5.678 MeV and sigma of 0.562 MeV. The amplitude of this Gaussian is left as a free parameter, and added to the predictive model in $\bar{\nu}_e$ energy for the PROSPECT measurement. By minimizing

the global χ^2 value, an amplitude of the Daya Bay excess is fit to the PROSPECT data. By definition of the fit parameter, an amplitude of 1.0 corresponds to an excess of the same scale as measured by Daya Bay, and an amplitude of 0.0 corresponds to no excess. For the PROSPECT data, an amplitude of 0.84 ± 0.39 is observed. This yields a simply-modeled excess in PROSPECT consistent with the excess measured by Daya Bay which disfavors a null-hypothesis of no excess at 2.17σ. Additionally, because of the average ^{235}U fission fraction at Daya Bay of 56% and the greater than 99% ^{235}U fission fraction at PROSPECT, an amplitude of 1.78 would correspond to the hypothesis where ^{235}U is the sole source of the model discrepancy in this energy range around 5.678 MeV in $\bar{\nu}_e$ energy. This hypothesis is also disfavored by the PROSPECT measurement at 2.44σ.

Chapter 5

Joint Spectral Analysis with Daya Bay

In addition to the results of the PROSPECT experiment, this thesis describes the joint spectral analysis between the PROSPECT and Daya Bay measurements. These experiments both measure the ^{235}U antineutrino prompt spectrum using complementary experimental techniques, while the Daya Bay experiment also deconvolves the ^{239}Pu prompt spectrum from its total LEU measurement . While it has been referenced previously in this manuscript, a more thorough discussion of the Daya Bay detector and analysis, focusing on similarities and differences between PROSPECT and Daya Bay, is included. Following that is a detailed discussion of the joint analysis between the two measurements.

5.1 Daya Bay

The Daya Bay experiments took data from December of 2011 until the same month of 2020 measuring antineutrinos produced by six commercial low enriched uranium (LEU) reactors in the Daya Bay nuclear power complex. The main goal of the experiment was to make a precision measurement of θ_{13}, the last 3-flavor neutrino oscillation parameter to be measured. The detector consisted of eight, identically designed monolithic antineutrino detectors (ADs). Four ADs were located in two near halls ($\sim$ 500m baseline from the

reactors) and four ADs were located in a far hall ($\sim$ 1600m baseline). Each AD consisted of three nested cylindrical vessels. The inner acrylic vessel contained 20 tons of liquid scintillator doped with 0.1% Gd as a neutron capture agent and acted as the main target volume of the AD. Surrounding that was another acrylic vessel with 20 tons of undoped liquid scintillator used to measure any energy escaping from the main target volume. Finally outside of that was a stainless steel containment vessel containing 192 PMTs facing the inner volume coupled via 37 tons of mineral oil that doubled as neutron shielding.

Daya Bay identified reactor antineutrino events using the same interaction as PROSPECT, inverse beta decay (IBD). Antineutrinos that interacted via IBD promptly deposited energy from the positron in the scintillator, while neutrons that thermalized in the detector could capture on the Gd, resulting in a gamma cascade of $\sim$8 MeV of energy in the delayed signal.

By looking at the ratio of neutrino spectra at two baselines, a relative $\bar{\nu}_e$ disappearance could be measured and a mixing angle could be quantified. The results from [23] are presented in Figure 5.1 using 1958 days of operational data. The left plot gives the total LEU spectrum at both near and far baselines, with a panel showing the ratio of the two spectra and the best fit oscillation signal. The right plot shows the 1, 2, and 3σ contours for the mixing angle and mass squared splitting terms for the fit, with best fit values of $\sin^2 2\theta_{13} = 0.0856 \pm 0.0029$ and $\Delta m^2_{32} = (2.471^{+0.068}_{-0.070}) \times 10^{-3}$ eV^2 (Δm^2_{32} assumes Normal Mass Heirarchy).

In addition to the measurement of the mixing angle, Daya Bay also deconvolves the full LEU spectrum into individual contributions from ^{235}U and ^{239}Pu. As individual reactor cores burn fuel, their relative fraction of fissile isotopes changes over time. A total effective fission fraction f^{eff} can be calculated as a function of time t. For each isotope i observed by detector d, it can be calculated weighing the fission fraction from each reactor r by its thermal power W, reactor-to-detector baseline L, and energy per isotope fission e as

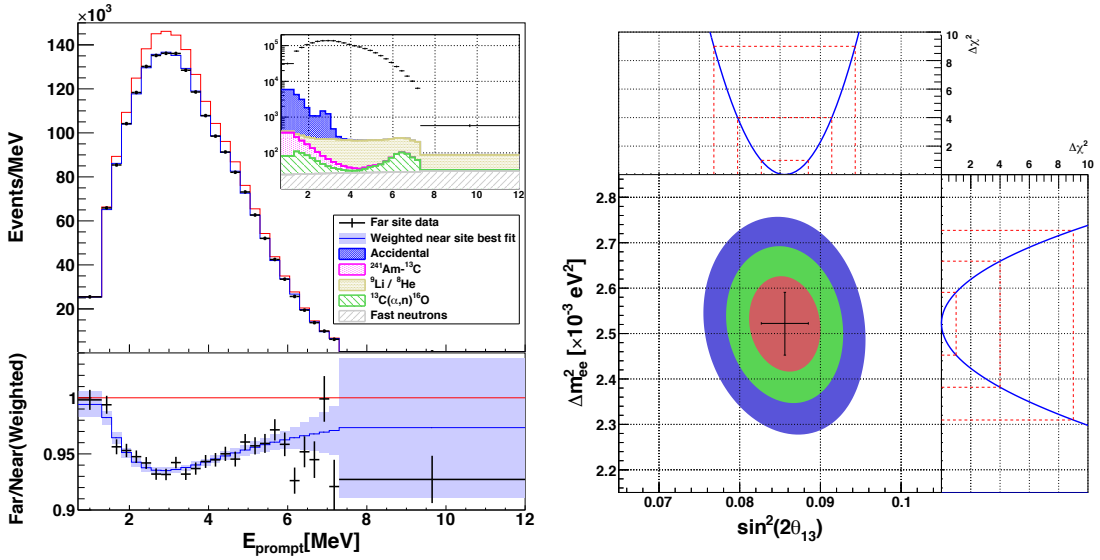

Figure 5.1: (left) Near and far hall measured antineutrino spectra from 1958 days of operation at Daya Bay. A ratio of the two spectra, as well as the best fit oscillation signal is shown in the lower plot. (right) The 1, 2, and 3σ contours for the mass squared splitting and mixing angle fit parameters [23].

follows:

$$f_{id}^{\text{eff}}(t) = \sum_r \frac{W_r(t) f_{ir}(t)}{L_{rd}^2 \Sigma_j f_{jr}(t) e_j} / \sum_r \frac{W_r(t)}{L_{rd}^2 \Sigma_j f_{jr}(t) e_j} \tag{5.1}$$

The fission fraction of each of the four main fissile isotopes (^{235}U, ^{238}U, ^{239}Pu, and ^{241}Pu) is calculated for the full Daya Bay dataset on a weekly basis. The left panel of Figure 5.2 [83] shows the total effective fission fraction of the other three isotopes as a function of the fission fraction of ^{239}Pu. The negative correlation between ^{235}U and ^{239}Pu is due to the nature that the burn-up of ^{235}U feeds into the production of ^{239}Pu in the reactor. In order to extract the individual contributions of ^{235}U and ^{239}Pu to the total LEU spectrum, a χ^2 is minimized defined as follows:

$$\chi^2(\eta^5, \eta^9) = 2 \sum_{djk} (S_{djk} - M_{djk} + M_{djk} \ln \frac{M_{djk}}{S_{djk}}) + f(\epsilon, \Sigma) \tag{5.2}$$

In this definition, d is the detector index, j is the index of one of 20 groups of data binned in ^{239}Pu fission fraction, k is the energy bin in prompt energy, M is the measured

prompt energy spectrum, ϵ is the set of nuisance parameters, $f(\epsilon, \Sigma)$ is a constraint term for nuisance parameters with correlations Σ and

$$S_{djk} = \alpha_k(\epsilon)s_k^5(\eta_k^5) + \beta_k(\epsilon)s_k^9(\eta_k^9) + s_k^{238+241}(\epsilon) + c_k(\epsilon) \tag{5.3}$$

Here, the expected prompt spectrum S is calculated using with $s^5(\eta^5)$ [$s^9(\eta^9)$] is the element of extracted ^{235}U (^{239}Pu) spectrum, α [β] is the corresponding coefficient for ^{235}U (^{239}Pu) factoring in detector mass, efficiency, baseline, and number of fissions where $s^{238+241}$ is the total predicted spectrum of ^{238}U and ^{241}Pu based off of the Huber and Mueller model with an extra 15% (10%) assigned uncertainty. η^5 and η^9 are the parameters fit over this χ^2 minimization to determine the total, deconvolved spectra for ^{235}U and ^{239}Pu respectively. Those final spectra are reported in the right panel of Figure 5.2, which also include a comparison to a scaled Huber model to explore a shape-only comparison. The local deviation of the measurement from the model reaches 4σ around 5 MeV in prompt energy.

5.2 Unfolding

The measured energy space of the PROSPECT and Daya Bay experiments reflect a variety of different systematic effects. Because Daya Bay's prompt energy space is very nearly measuring the energy of the positron and annihilation gammas from the IBD while PROSPECT's reconstructed visible energy incorporates effects from the scintillator non-linearity and a non-trivial effect from escaping energy, the spectra cannot be compared directly. This effect is shown in Figure 5.3 [84], where the top panel shows the response of the PROSPECT and Daya Bay detectors on the detected signal from neutrinos at three different energies in the top panel. The relative shift in the peaks of the distributions is predominantly from the non-linearity of the scintillator, which is calibrated out of the Daya

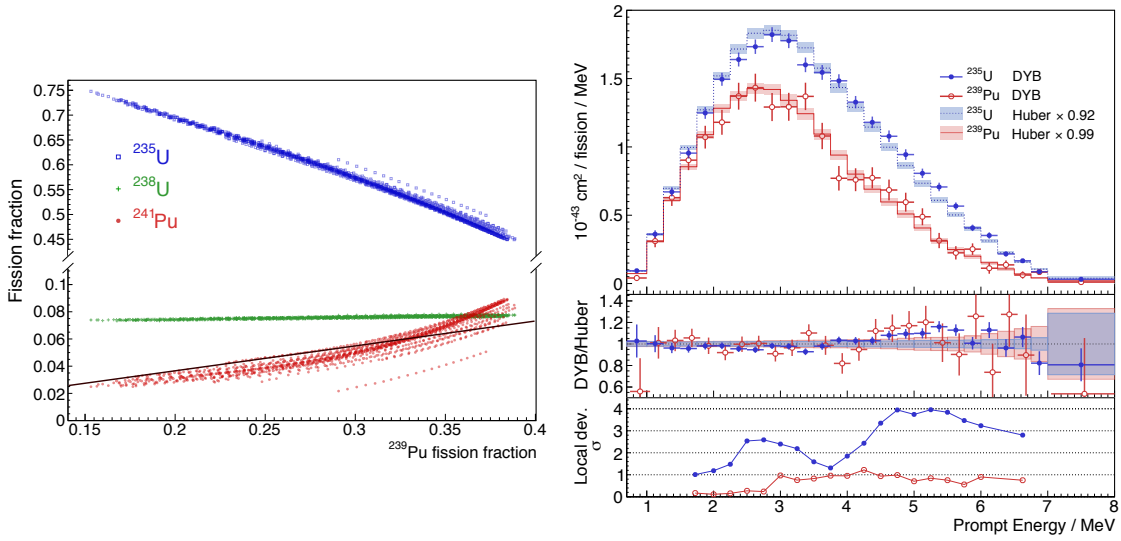

Figure 5.2: (left) Total effective fission fraction of the fissile isotopes as a function of the fission fraction of ^{239}Pu. A negative correlation is evident with respect to the fission fraction of ^{235}U which feeds into ^{239}Pu as it burns up. (right) The final fitted spectra of ^{235}U and ^{239}Pu as deconvolved from the full LEU measurement. the deviation from scaled Huber models are given for comparison, as well as the local deviations of the data from the model [83].

Bay signal but incorporated into the PROSPECT measurement. The difference in widths of the peaks is predominantly from systematics associated with inactive detector volume in the PROSPECT detector. The effect of the latter is quantified in the lower panel, where the full width of the peak at half the maximum value of the distribution is plotted as a function of energy.

Due to this, the measurements must be transformed into an identical energy space in order to make a meaningful comparison. This is handled in two ways: one by unfolding each measurement into antineutrino energy and two by refolding the measurement from one experiment into the energy space of the other.

For the first compatibility measurement technique, each set of experimental data can be unfolded from its respective prompt energy space into antineutrino energy space. For this, the Wiener-SVD [85] method is used. This method is used as a means to determine a regularization in order to optimize the tradeoff between the relative uncertainty of fit pa-

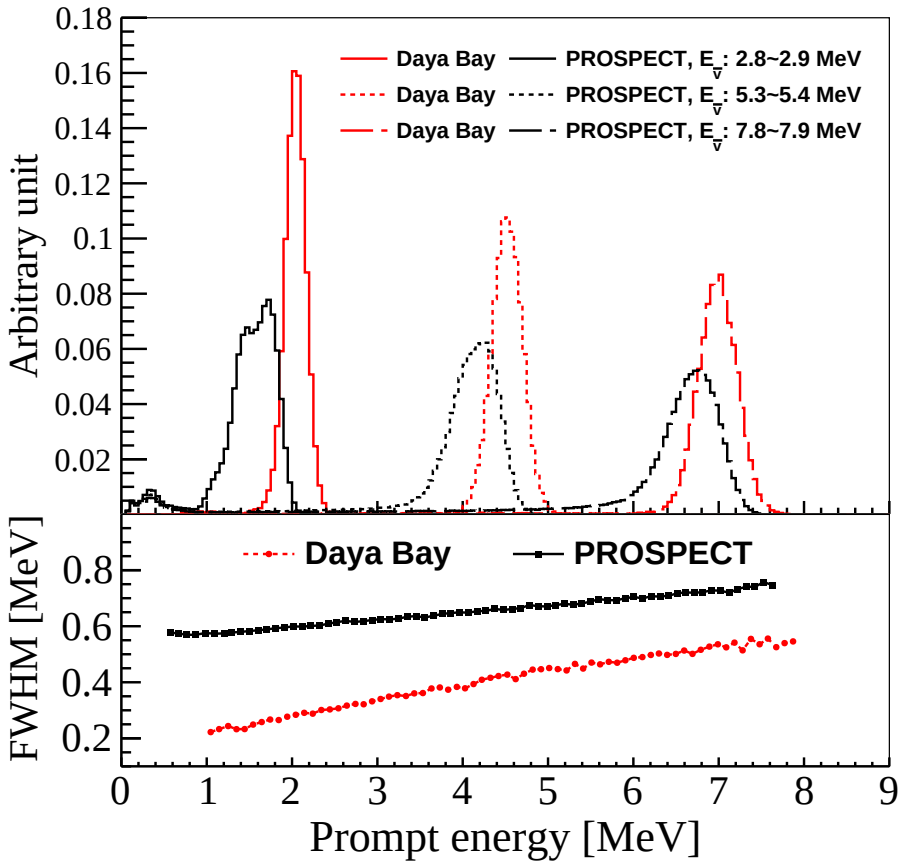

Figure 5.3: (Top) Reconstructed prompt energy distributions based on the $\bar{\nu}_e$ signals with specific energy ranges (uniform distribution). The distributions are normalized to 1. The shift in peak location is driven primarily by different effects from scintillator non-linearity in the energy response. The difference in FWHM is primarily due to different effects from inactive volume in the detector response functions. (Bottom) FWHM versus prompt energy corresponding to the peak for the reconstructed prompt energy distributions from specific energy ranges [84].

rameters and introduced bias. This is important due to the nature of the unfolding problem as follows. Because there are real statistical fluctuations or noise on top of the neutrino signal, any places where the scale of the noise is comparable to or larger than the signal, the inverted spectrum (in neutrino energy) is prone to significant effects from amplified statistical fluctuations. For example, if $\vec{S}$ is the true energy spectrum of a neutrino signal, R is the response matrix of the detector with R^{-1} its inverse (assuming R is invertible), $\vec{M}$ is the measured signal, and $\vec{N}$ is the noise component the measurement, then one can calculate the estimated true energy signal $\tilde{S}$ as

$$\begin{aligned} \tilde{S} &= R^{-1} \cdot \vec{M} \\ &= R^{-1} \cdot (R \cdot \vec{S} + \vec{N}) \\ \tilde{S} &= \vec{S} + R^{-1} \cdot \vec{N} \end{aligned} \tag{5.4}$$

where $R^{-1} \cdot \vec{N}$ results in an error in the estimate of $\vec{S}$. In general, this can be addressed by implementing a regularization term into the unfolding procedure to limit the effect of noise in areas of small signal, which limits the large fluctuations in the uncertainty of the unfolded spectrum. Additionally, the response matrix R and the measured data M can be "pre-whitened" by left multiplying by Q, the lower triangular matrix of the Cholesky decomposition of the inverse of the Covariance matrix. This undoes some of the correlated uncertainties in the measurement from the response matrix and the data, and all uses of R and M in this section are treated as "pre-whitened". By expressing R^{-1} as decomposed by Singular Value Decomposition, one can implement a regularization filter F. Let D be a (possibly non-square) matrix where D_{ii} is the ith singular value d_i in decreasing value of R and $D_{ij} = 0$ for $i \neq j$. Also let $D^{-1}_{ii} = \frac{1}{d_i}$ with $D^{-1}_{ij} = 0$ for $i \neq j$. Then

$$\begin{aligned} R &= UDV^T \\ \tilde{R}^{-1} &\coloneqq VFD^{-1}U^T \end{aligned} \tag{5.5}$$

Where $F_{ii} = \frac{d_i^2}{d_i^2+\tau}$ where τ is a regularization term. The implementation of this regularization term, while smoothing out the effect from unfolding noise leads to biases in the estimate $\tilde{S}$. One can see in the filter term F that as the regularization term τ goes to zero, F goes to an identity operator, leaving D^{-1} unchanged, whereas increasing τ dampens out effects of small d_i. This is the mechanism for suppressing the large fluctuations when noise in the measurement is not negligible compared to the signal. By tuning τ, one can dampen out more of the noise's effects in the unfolded signal estimate at the expense of adding more bias into the estimate. Using the Wiener filter, one can optimize the balance between this tradeoff in maximizing noise suppression and suppressing bias. This form of the Wiener filter is derived from signal processing, where the filter takes the form of the true signal squared over the sum of the squares of the true signal and an additive noise term squared. Specifically, the filter W (to be swapped in for F in the previous equation) is defined as

$$W_{ii} = \frac{d_{Ci}^2 \cdot \left(\Sigma_j V_{Cij}^T \cdot (\Sigma_l C_{jl} \cdot \overline{s}_l)\right)^2}{d_{Ci}^2 \cdot \left(\Sigma_j V_{Cij}^T \cdot (\Sigma_l C_{jl} \cdot \overline{s}_l)\right)^2 + 1} \tag{5.6}$$

where C is a second derivative assisting matrix

$$C = \begin{bmatrix} -1+\epsilon & 1 & 0 & \dots & 0 & 0 & 0 \\ 1 & -2+\epsilon & 1 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 1 & -2+\epsilon & 1 \\ 0 & 0 & 0 & \dots & 0 & 1 & -1+\epsilon \end{bmatrix} \tag{5.7}$$

with $\epsilon = 10^{-6}$, $\overline{s}$ is the expectation of the true spectrum (the Huber model in this case), and d_C and V_C are defined by

$$R \cdot C^{-1} = U_C \cdot D_C \cdot V_C^T. \tag{5.8}$$

Ultimately, this can all be combined to give an estimate of the unfolded spectrum $\tilde{S}$

$$\tilde{S} = C^{-1} \cdot V_C \cdot W \cdot V_C^T \cdot C \cdot (R^T \cdot R)^{-1} \cdot R^T \cdot M \tag{5.9}$$

which optimizes the trade off between statistical uncertainty and introduced bias the unfolded spectrum.

5.3 Compatibility

The compatibility of the two measurements is done using a χ^2 test in neutrino energy space using covariance matrices in three separate ways. One way is directly comparing the two unfolded spectra using a covariance matrix generated under the assumption that both experiments are based on the same underlying Huber spectrum. Compatibility is also done by jointly fitting a spectrum to the two measurements and evaluating that fitted χ^2 with the relevant degrees of freedom. Finally, the spectra are compared by refolding the Daya Bay prompt measurement into the reconstructed energy space of PROSPECT. These three methods will be referred to as the direct comparison, the fitted comparison, and the refolded comparison respectively in the following subsections.

5.3.1 Direct Comparison Method

Using the two unfolded measurements from PROSPECT and Daya Bay, $\tilde{S}_P$ and $\tilde{S}_D$, the direct comparison figure of merit is defined as follows:

$$\chi^2_{DC} = (a\tilde{S}_P - \tilde{S}_D)^T \, V_{DC}^{-1} \, (a\tilde{S}_P - \tilde{S}_D) \tag{5.10}$$

where a is a floating normalization parameter scaling $\tilde{S}_D$ over which χ^2_{DC} is minimized (since there is no rate absolute rate information between $\tilde{S}_P$ and $\tilde{S}_D$) and V_{DC} is computed

numerically using ten thousand toys for each experiment. The toys are generated from a model-based prediction in prompt energy space with fluctuations based on real experimental uncertainties (including correlated and uncorrelated contributions). These toys are then unfolded into neutrino energy space using the Wiener-SVD technique, and a covariance matrix is calculated using the differences of the unfolded toys from the truth models. Additionally, for the covariance matrix generation an extra uncertainty is incorporated into the unfolding process. A random 3% uncorrelated fluctuation is added to each bin of the model $\overline{s}$ (Eq. 5.6) for each pair of toys used to generate the covariance matrix. This incorporates an uncertainty into the prior used to optimize the Wiener filter and reduces the explicit model dependency of the analysis. The final covariance matrix is made by averaging over these ten thousand toy covariance matrices. Elements of the direct comparison covariance matrix for N toy based neutrino energy spectrum $\tilde{S}$ can be defined as:

$$V_{DC}(i,j) = \frac{1}{N}\sum_{n=1}^{N}\left[\left(f\tilde{S}_{Pn}(i) - \tilde{S}_{Dn}(i)\right) \times \left(f\tilde{S}_{Pn}(j) - \tilde{S}_{Dn}(j)\right)\right] \tag{5.11}$$

where f is a scaling factor setting the integral of the PROSPECT spectrum equal to the integral of the Daya Bay spectrum. Since the toys are generated using the uncertainties from the real covariance matrices of each experiment, and since those toys are unfolded into neutrino energy, this covariance matrix contains uncertainties both involving the subtraction of spectra with original uncertainties unfolded into neutrino energy as well as contributions from the bias of the unfolding. It is presented in Fig 5.4 for reference.

As a means of evaluating the sensitivity of this comparison, one thousand pairs of toys were compared. Using a sliding 1.5 MeV window, the local deviation between the toys was evaluated by adding free floating parameters to 6 consecutive bins (.25 MeV each) as well as an overall normalization parameter. χ^2_{DC} is then minimized over all of those parameters, and the difference is taken between that result and the χ^2_{DC} when there are no extra fit parameters besides the initial scaling factor a used in equation 5.10. A local

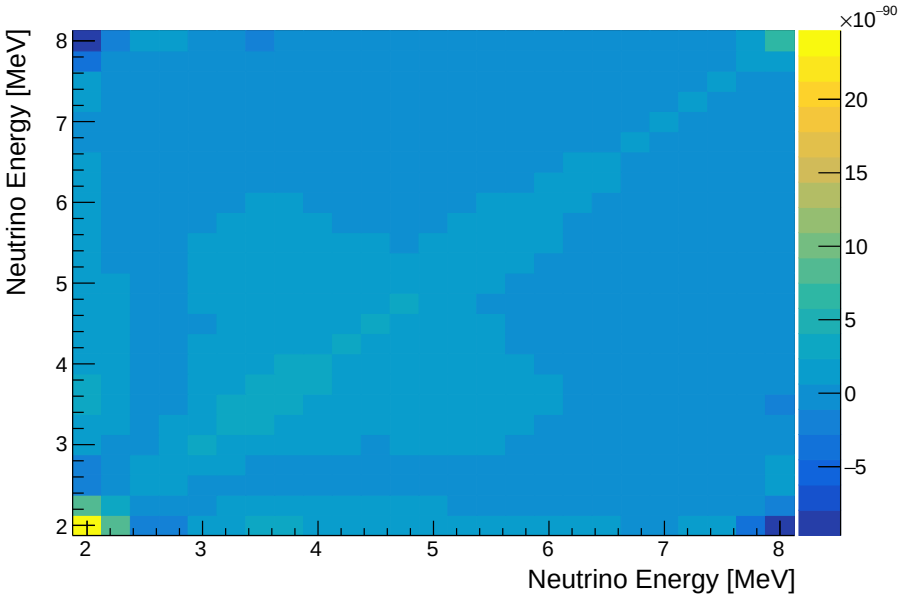

Figure 5.4: Covariance matrix made from the differences between ten thousand pairs of toys unfolded from prompt energy into neutrino energy for both the PROSPECT and Daya Bay experiments.

p-value is calculated based on this difference relative to a χ^2 distribution with six degrees of freedom. For each sliding energy window, the median local p-value is found for one thousand toys, and plotted in Figure 5.5 for a comparison of toys with identical input models. This process is then repeated for cases where the input PROSPECT and Daya Bay models differ by the addition of a gaussian bump with a mean of 5.68 MeV and a width of 0.57 MeV, values roughly equal to the bump seen in previous θ_{13} measurements like Daya Bay. One set of models had an amplitude of the gaussian equal to 10% of the value of Huber at the mean of the bump (approximately equal to the deviations seen in the full Daya Bay spectral measurement), and another set of models had the addition of a gaussian with an amplitude of 20% that of Huber. The dashed horizontal lines are the one and two sigma thresholds. This figure shows the sensitivity of this particular analysis to differences between the measurements. It showcases that even if Daya Bay's measurement had an LEU-sized bump while PROSPECT's measurement had no bump, this test would not typically show more than a 1σ deviation.

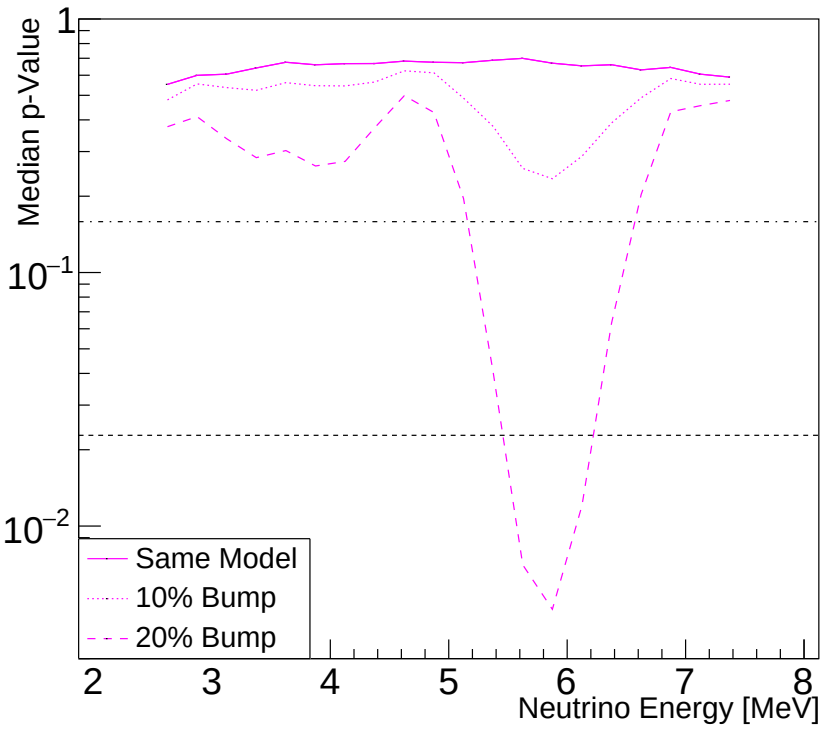

Figure 5.5: Median local p-value contribution across 1.5 MeV for the compatibility of two toys generated by identical models, models differing by a Daya Bay like bump (10% of the Huber spectrum at the mean), and twice that (20% of the Huber spectrum at the mean).

Because of the introduction of the additional 3% uncertainty on the input model in the unfolding process, the χ^2 does not follow the analytical distribution for 24 degrees of freedom. To find the proper p-value, a histogram generated from 1000 Huber-based toys is made, and the p-value is calculated numerically from the histogram. This numerically generated distribution is shown in Figure 5.6 both with the inclusion of the 3% unfolding uncertainty (which is actually used to calculate the p-value) as well as a distribution generated without the extra 3% uncertainty as a crosscheck that it reproduces the expected analytical distribution.

For the actual measurement comparison, the scaling of the PROSPECT measurement to the Daya Bay measurement is done with a freely floating fit parameter minimizing the χ^2. Using this method, the measurement yields a χ^2/dof of 13.2 / 24 corresponding to a p-value of 0.90. A breakdown the contributions toward this measurement as a function of energy are included in Figure 5.7. The local p-value is calculated by taking the difference of the minimized χ^2 and a χ^2 calculated with the six additional degrees of freedom from additional fit parameters on adjacent bins (similar to the sliding window analysis described in Section 4.3.3. This shows the contribution towards the global χ^2 of a 1.5 MeV wide sliding window.

5.3.2 Jointly Fit Spectrum

The second method of comparing the two measurements is done by doing a joint fit to a model, and taking the $\Delta\chi^2$ when fitting each measurement to the model independently. Here, the test statistic is defined as follows:

$$\Delta\chi^2 = \chi^2_{jnt} - \chi^2_{ind} \tag{5.12}$$

where

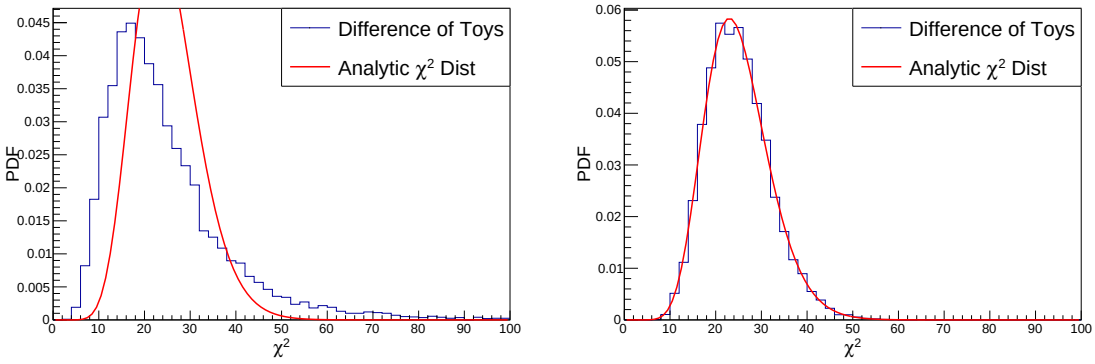

Figure 5.6: χ^2 distributions generated using 1000 pairs of Huber-based toys compared to the expected analytical distribution. These toys are evaluated using equation 5.10. The left plot includes an additional 3% uncertainty on the model used in the Wiener filter from unfolding while the right plot omits that uncertainty as a crosscheck to reproduce the expected analytical distribution. The left plot is used to determine the real p-value for the data.

$$\chi^2_{ind} = \Delta(\tilde{S}_P, \overline{s}, \alpha)^T V^{-1}_{JF,P} \Delta(\tilde{S}_P, \overline{s}, \alpha) + \Delta(\tilde{S}_D, \overline{s}, \beta)^T V^{-1}_{JF,D} \Delta(\tilde{S}_D, \overline{s}, \beta) \tag{5.13}$$

$$\chi^2_{jnt} = \Delta(\alpha_0 \tilde{S}_P, \overline{s}, \alpha)^T V^{-1}_{JF,P} \Delta(\alpha_0 \tilde{S}_P, \overline{s}, \alpha) + \Delta(\tilde{S}_D, \overline{s}, \alpha)^T V^{-1}_{JF,D} \Delta(\tilde{S}_D, \overline{s}, \alpha) \tag{5.14}$$

and $\Delta(\tilde{S}_P, \overline{s}, \alpha)$ is the difference between the unfolded measurement $\tilde{S}_P$ and the Huber model $\overline{s}$ with a vector of freely floating fit parameters α_i for each energy bin ($i = 1, ..., 25$), including an additional relative normalization parameter α_0 for use in the PROSPECT measurement for χ^2_{jnt}, having no absolute rate information. The only difference between the definitions of χ^2_{jnt} and χ^2_{ind} is that the fit parameters are independent of each other in χ^2_{ind} and jointly constrained in χ^2_{jnt}. In practice, χ^2_{ind} is approximately zero up to numerical precision, and the value for $\Delta\chi^2$ is simply χ^2_{jnt}. The full definition is provided for clarity and formalism for determining the appropriate number of degrees of freedom.

The covariance matrices are generated similarly to the way done in the Direct Comparison method in Eq 5.11. Instead of taking the difference between toys of the different

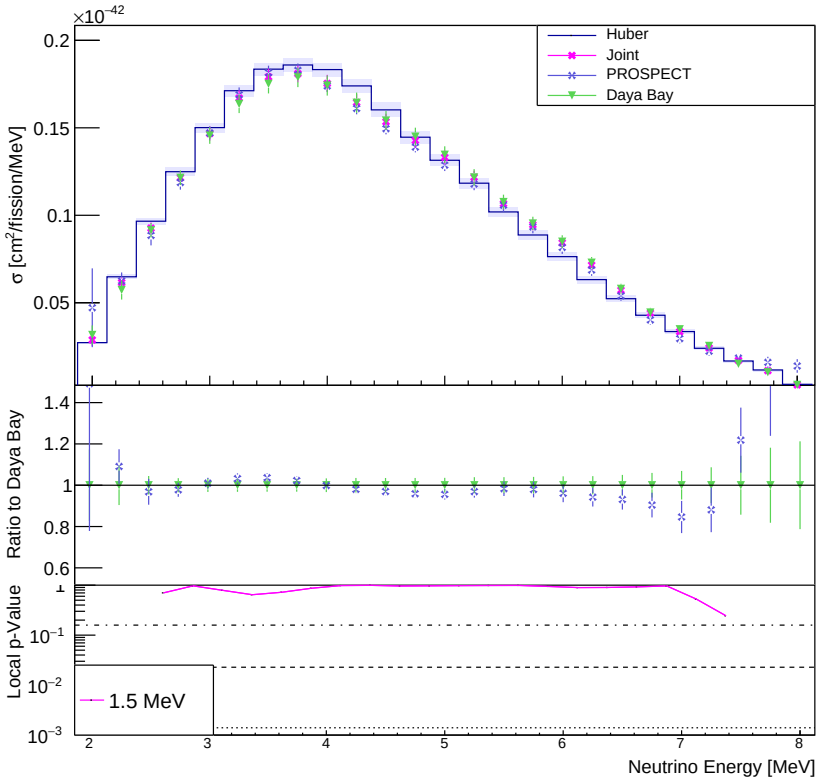

Figure 5.7: (top) Unfolded spectra for both PROSPECT and Daya Bay with Huber plotted as a reference. (middle) Ratio of PROSPECT spectrum to Daya Bay with square root of the diagonals of the each experiment's covariance matrix used for error bars. (bottom) Local p-value of the discrepancy between the two measurements using a 1.5 MeV wide energy window.

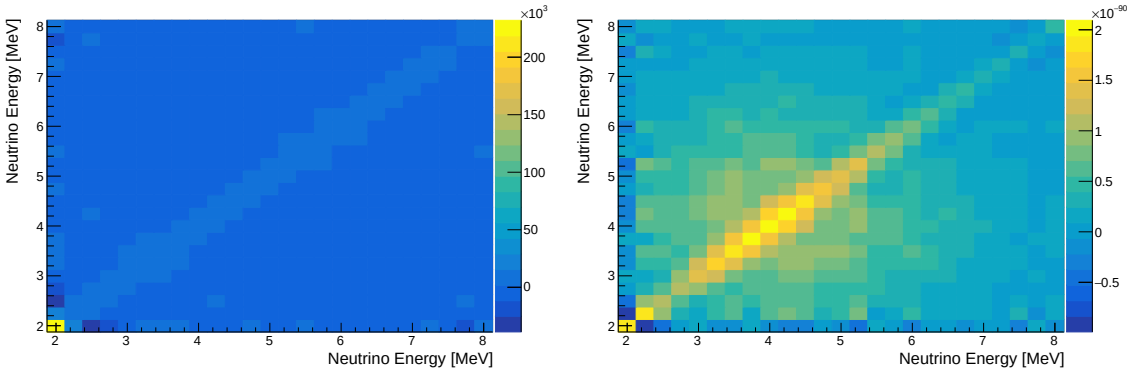

Figure 5.8: Covariance matrices for PROSPECT (left) and Daya Bay (right) in unfolded neutrino energy space.

measurements, they are generated taking the difference from prompt toys unfolded into neutrino energy space and the Huber model used to produce them. This model is used here since this is the model used to optimize the unfolding process in the previous section. The model is scaled to the toys by area here.

$$V_{JF}(i,j) = \frac{1}{N}\sum_{n=1}^{N}\left[\left(\tilde{S}_{Pn}(i) - \overline{s}(i)\right) \times \left(\tilde{S}_{Pn}(j) - \overline{s}(j)\right)\right] \tag{5.15}$$

The joint fit covariance matrices V_{JF} are constructed using $N = 10,000$ toys for each experiment. They are shown in Fig 5.8 for reference.

For measurements in neutrino energy space with 25 bins, this leads to χ^2_{ind} having 0 degrees of freedom (50 bins with 50 independent fit parameters), and χ^2_{jnt} having 24 degrees of freedom (50 bins with 26 independent fit parameters). Thus, $\Delta\chi^2$ has 24 degrees of freedom, noting that χ^2_{ind} is effectively zero within machine precision. After performing the fit itself, this yields a χ^2/dof of 16.8 / 24 corresponding to a p-value of 0.85.

5.3.3 Refolding Method

The final comparison method transforms one prompt measurement directly into the prompt energy space of the other measurement. In order to avoid the amplification of statistical

features, it is crucial that the prompt energy space with less overall energy smearing is refolded into the prompt energy space with more effective energy smearing. Looking at the bottom panel of Figure 5.3 [84], it is clear that refolding the Daya Bay prompt energy space into the PROSPECT energy space will alleviate this potential issue.

The transformation of the Daya Bay prompt measurement M_p^{DYB} to the newly mapped spectrum M_{map}^{DYB} in the PROSPECT energy space is done via a mapping matrix R^{map}

$$M_{map}^{DYB} = R^{map} M_p^{DYB} = R^{PRO}(R^{DYB})^{-1} M_p^{DYB}. \tag{5.16}$$

Note that since R^{DYB} is not a square matrix, $(R^{DYB})^{-1}$ is really a pseudo inverse similarly defined in equation 5.5. The comparison of the two data sets in PROSPECT reconstructed energy is included in Figure 5.9 [84] where there is an additional factor used to scale the total integrated flux to the absolute rate as measured by Daya Bay.

By similarly transforming the Daya Bay covariance matrix V_p^{DYB} to V_{map}^{DYB}, a new covariance matrix for the Daya Bay measurement in the PROSPECT energy space is calculated:

$$V_{map}^{DYB} = R^{map} \times V_p^{DYB} \times (R^{map})^T, \tag{5.17}$$

and the square root of the diagonals of this matrix are used for the error bars in Figure 5.9 [84].

Then a test statistic χ^2_{refold} is defined to quantitatively evaluate the compatibility of the measurements. Using the fitting approach first described in Equation 5.14 in Section 5.3.2, and modifying it to work with the mapping matrix R^{map} one gets

$$\begin{aligned}\chi^2_{refold} =& (S^{fit} - M_p^{DYB})^T (V^{DYB})^{-1} (S^{fit} - M_p^{DYB}) \\ &+ (\alpha_0 R^{map} S^{fit} - M_p^{PRO})^T (V^{PRO})^{-1} (\alpha_0 R^{map} S^{fit} - M_p^{PRO}).\end{aligned} \tag{5.18}$$

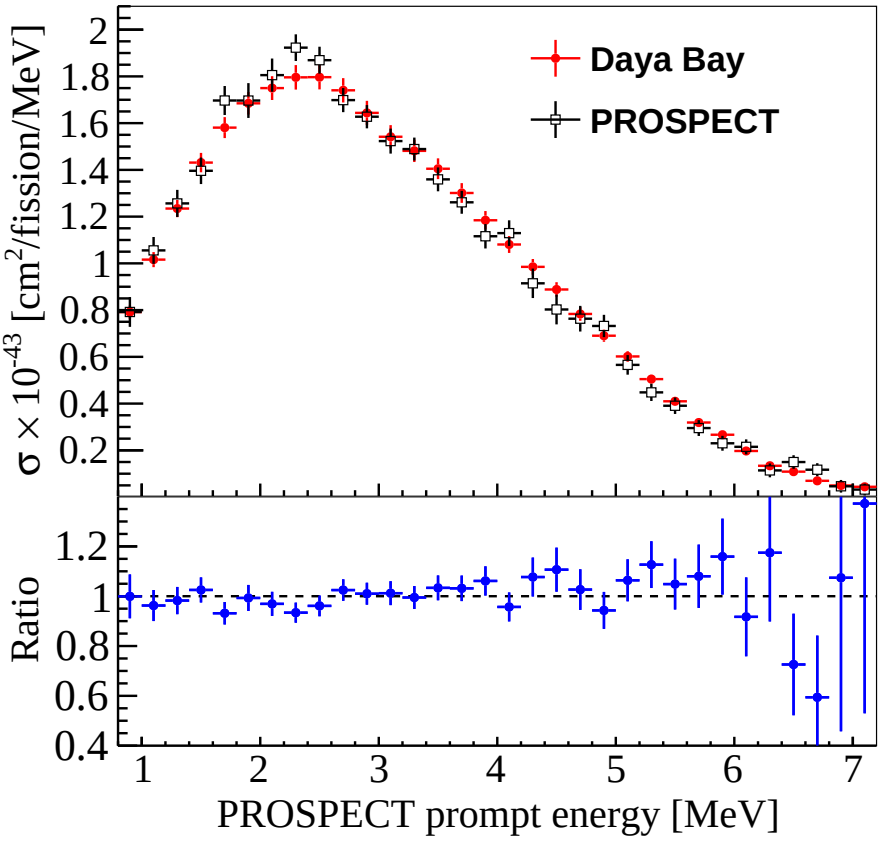

Figure 5.9: (top) Comparison of the two ^{235}U measurements in the reconstructed visible energy of PROSPECT (bottom) and the ratio of the Daya Bay measurement to PROSPECT. All Daya Bay data has been refolded using both response matrices, and the y-axis has been scaled to match the total integrated flux as measured by Daya Bay. The two measurements are consistent across all energies [84].

Here S^{fit} is a fitted spectrum with independent, freely floating parameters on each of the 26 energy bins corresponding to the Daya Bay prompt energy spectrum. This spectrum is then mapped to the PROSPECT prompt energy space, and allowed to globally renormalize via the fit parameter α_0. Just like in Section 5.3.2, χ^2_{refold} can be compared to a freely floating fit on just the Daya Bay measurement ($\chi^2 = 0$, by construction) such that the meaningful $\Delta\chi^2$ test statistic is equal to χ^2_{refold}. Thus comparing this value to an expected χ^2 distribution with 31 degrees of freedom (32 additional bins and 1 new fit parameter for the $\Delta\chi^2$ metric) can find the compatibility of the two measurements without the need to introduce any additional smearing due to the unfolding procedure. Specifically, the data gives a χ^2/dof of 25.44/31, corresponding to a p-value of 0.75, again showing the data are consistent. This p-value is also the most sensitive of the three methods due to this removal of smearing from the unfolding technique.

This result shows general agreement with the Direct Comparison Method and Jointly Fit Method. It should be noted that the Direct Comparison Method incorporates any correlated effects from unfolding that are present in both measurements. While this effect is small, it is ignored in the Joint Fit method which treats the two measurements as completely independent. Ultimately, though, the refolding metric eliminates any prominent effects of the smearing introduced by unfolding into the data comparison, and is the main result of the comparison analysis. However, all three results show no significant discrepancies between the Daya Bay and PROSPECT ^{235}U measurements.

5.4 Joint Spectral Fit

With both experiments showing consistency with each other, their measurements can be combined to produce one jointly constrained measurement of the ^{235}U neutrino spectrum. In fact, one can generate a spectrum directly from the fit implemented in section 5.3.2. Taking the 25 energy parameters from the fit α and multiplying them by the starting spec-

trum values from $\overline{s}$, one can get the jointly constrained spectrum. The full covariance matrix for the measurement can be calculated as

$$C_{jnt} = D_{\overline{s}} C_{\alpha} D_{\overline{s}} \tag{5.19}$$

where $D_{\overline{s}}$ is a matrix with elements of $\overline{s}$ along the diagonal, and C_{α} is the covariance matrix of the α parameters from the fit defined in Section 5.3.2 ($\alpha \neq 0$). C_{α} is produced during the χ^2 minimization done by the ROOT fitter using the Minuit2 package [86], with uncertainties derived from the Hessian matrix of the fitted parameters (i.e. after calling the Hesse command). That fitted spectrum is shown in the top panel of Figure 5.10 along with the individually unfolded PROSPECT and Daya Bay spectra, as well as the Huber model for reference. The middle panel of the same figure shows the same spectra as a ratio to the Huber model, noting that the error bars plotted in both of these panels are the square root of the diagonals of the neutrino energy covariance matrices. Because of the highly correlated uncertainties in the transformed covariance matrices of Figure 5.8, simply plotting the errors this way may visually misrepresent the uncertainties. The bottom panel is the local deviation of the joint fit from the Huber model, calculated the same way as in Figure 5.7. This figure shows a local deviation from Huber of nearly 3σ around 6 MeV. The full spectrum comparison to Huber (including the Huber uncertainties) results in a χ^2/dof of 30.4 / 24, which corresponds to a p-value of 0.17. The full uncertainties on the fitted parameters are propagated to the final spectral bins including all correlations across each bin and the relative scaling factor as a full covariance matrix and included in Figure 5.11. A plot of the relative uncertainties of the PROSPECT, Daya Bay, and Joint spectra are given in Figure 5.12.

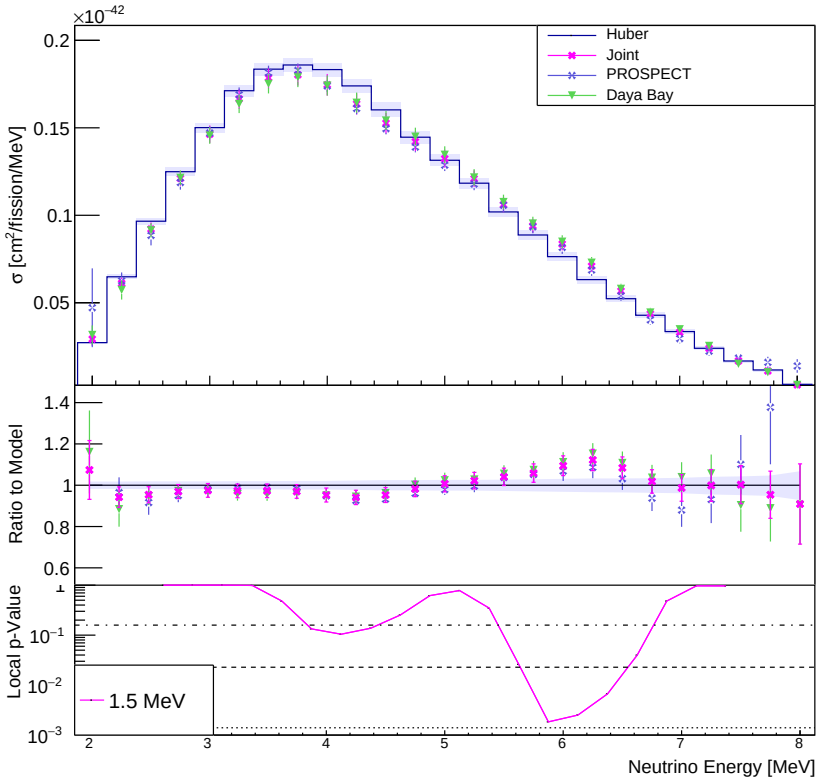

Figure 5.10: (Top) Unfolded ^{235}U spectrum from PROSPECT, Daya Bay, a joint fit constrained by both experiments, and the scaled Huber model for reference. All spectra are scaled to the integrated rate of the Daya Bay measurement. (Middle) Ratio of each of the three unfolded spectra to the scaled Huber model. (Bottom) Local p-value of the discrepancy between the joint fit and the Huber model using a 1.5 MeV wide energy window.

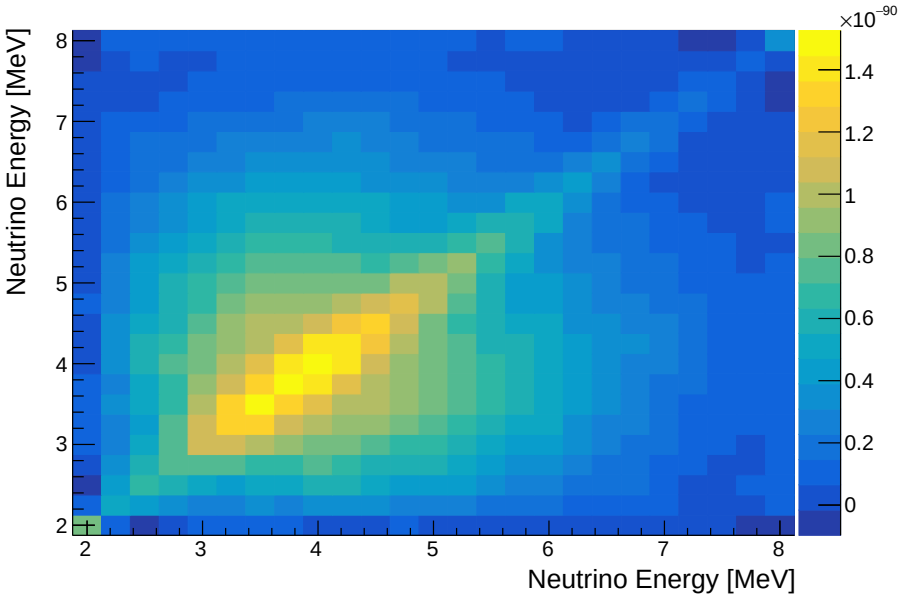

Figure 5.11: Full covariance matrix of the jointly fit spectrum in neutrino energy.

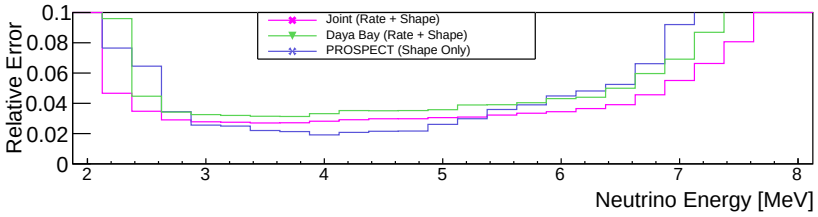

Figure 5.12: Relative uncertainties of the neutrino energy spectra for PROSPECT, Daya Bay, and the Joint fit. The PROSPECT measurement only has shape information, while the Daya Bay and Joint fit have absolute rate information in their uncertainties. This shows the improvement over Daya Bay alone that the PROSPECT shape information contributes to the joint fit. The relative uncertainties improve by roughly 0.5% across the entire spectrum (e.g. 3.5% goes to 3% relative uncertainty around 3 MeV)

	Daya Bay	PROSPECT	Joint
Spectrum	26×1	32×1	58×1
Response	26×25	32×25	58×25
Covariance	26×26	32×32	58×58

Table 5.1: Dimensions of individual input parameters for unfolding of the individual Daya Bay, individual PROSPECT, and joint analyses.

5.5 Jointly Unfolded Spectrum

In addition to the joint fit analysis, a jointly unfolded spectrum is also produced. This is only done since the two measurements are shown to be consistent with eachother. By doing a joint unfold, it allows for an integration of both datasets using a Wiener smearing matrix constrained by the combined signal to noise ratio of the measurements together rather than having separate smearing matrices as are calculated using the joint fit method. The joint inputs are combined using a direct sum of the individual inputs with dimensions as specified in Table 5.1. Block matrices and vectors showing the structure of the inputs are shown below, noting that all PROSPECT inputs are rescaled to have a matching absolute rate to Daya Bay in the unfolded space.

$$M_{JNT} = \begin{bmatrix} M_{DYB} \\ M_{PRO} \end{bmatrix}, \tag{5.20}$$

$$R_{JNT} = \begin{bmatrix} R_{DYB} \\ R_{PRO} \end{bmatrix}, \tag{5.21}$$

$$V_{JNT} = \begin{bmatrix} V_{DYB} & 0 \\ 0 & V_{PRO} \end{bmatrix}. \tag{5.22}$$

By using these new joint inputs in the same Wiener-SVD unfolding process as described in Section 5.2, one is able to produce a single, unfolded spectrum with a singular

smearing matrix A_C constrained by both experiments defined by

$$A_C = C^{-1} \cdot V_C \cdot W_C \cdot V_C^T \cdot C. \tag{5.23}$$

Here A_C is a smearing matrix that can be applied to an input model in order to directly apply the smearing introduced in the Wiener-SVD regularization directly. This effects the overall bin-to-bin correlations of the final measurements as demonstrated in Figure 5.13. This figure shows three slices of the covariance matrices produced with the joint fit and joint unfold methods. While the spectrum-shaped structure associated with the completely correlated uncertainty of the rate is very similar between the two methods, the features around 2.5, 4, and 6.5 MeV show a broader bin-to-bin uncertainty with the joint fit method. This represents a broader smearing introduced with that method, partially due to the fact that the PROSPECT measurement is a statistically-limited measurement. This means that an independent unfolding of PROSPECT will incorporate a greater smearing from the Wiener-SVD unfolding than a single unfolding with both the PROSPECT and Daya Bay measurements.

The final, jointly unfolded result is provided in Figure 5.14. This shows the unfolded results from the individual experiments as well as the jointly unfolded result using the joint Wiener-SVD defined above. The integrated rate for each analysis is constrained to equal that of the Daya Bay measurement, and the ratio to the Huber model is also shown. The ratio showcases both the general agreement of all analyses as well as the total flux deficit relative to the prediction. The bottom panel of the figure shows the total improvement of the final uncertainty over the Daya Bay-only result, while the shape-only PROSPECT relative uncertainty is also presented.

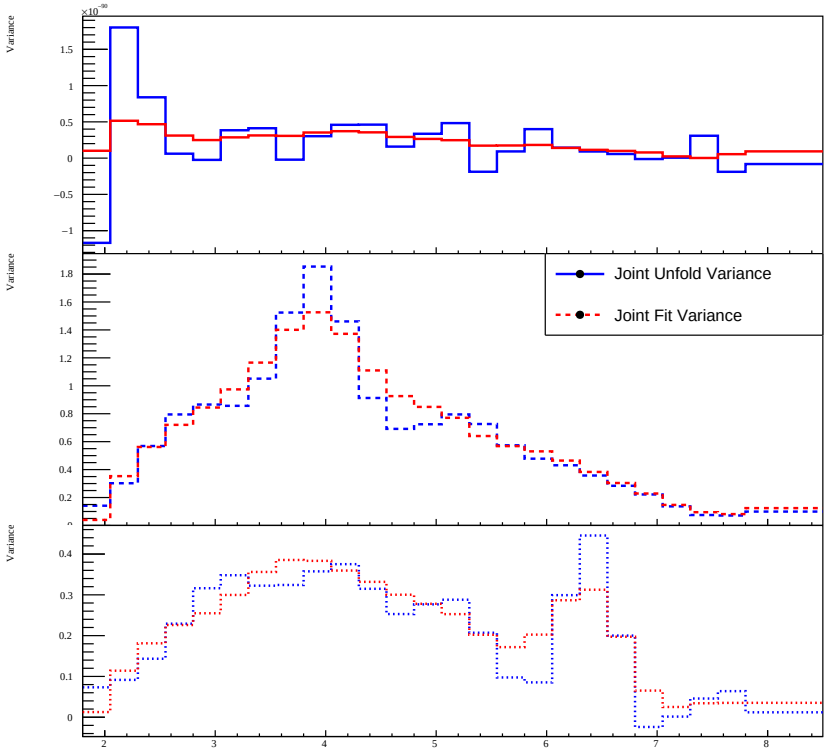

Figure 5.13: Slices of the covariance matrices obtained for the joint fit and joint unfold analyses. On top of the similar overall correlated uncertainties, the structures at $\sim$ 2.5, 4, and 6.5 MeV in the three panels show that the joint fit method has stronger bin-to-bin correlations associated with greater smearing.

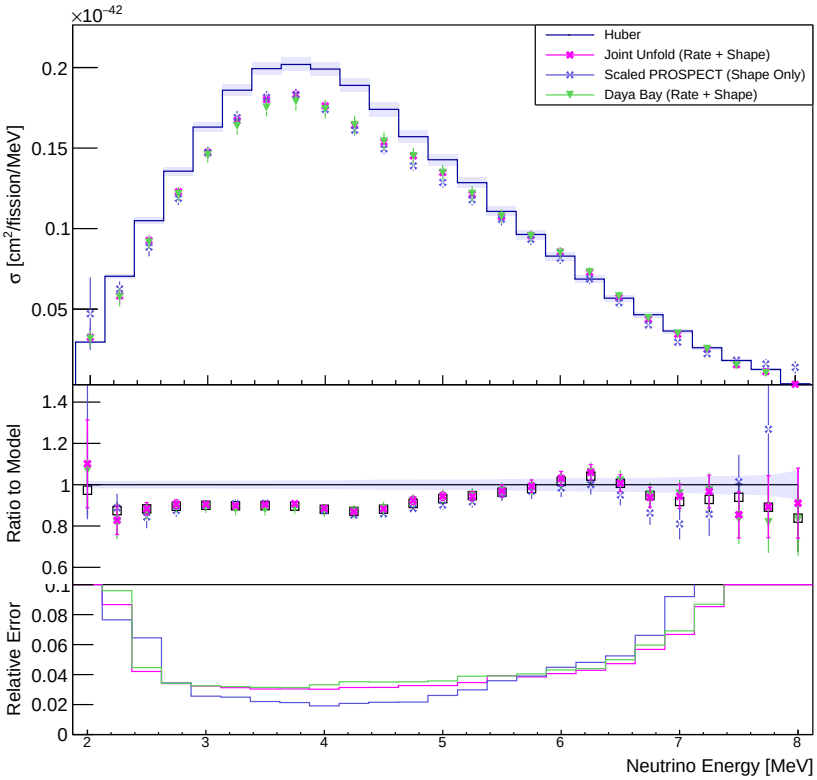

Figure 5.14: (top) Unfolding results from the individual unfolding of Daya Bay and PROSPECT as well as the jointly unfolded spectrum. All results are scaled to the absolute rate of Daya Bay, besides the Huber model which is shown for reference. (middle) ratio of each result to the model, showing both the strong agreement between results as well as the global deficit from Huber. (bottom) Relative error of each measurement showing the strict improvement of the joint analysis over the Daya Bay only rate and shape uncertainty via incorporation of the shape information from PROSPECT.

5.6 Joint Deconvolution

The final approach to for the joint analysis is to take the PROSPECT data and use it as an additional constraint in the prompt energy deconvolution of the Daya Bay LEU deconvolution. For the Daya Bay deconvolution, the analysis utilizes the different known fuel compositions of the various reactors measured to calculate a baseline-weighted effective fission fraction breakdown for the measurement as a function of time. This average fission fraction of ^{235}U ranges from approximately 0.5 to 0.7 across the whole dataset. By binning the measurement in fission fraction of ^{235}U (as shown in the left panel of Figure 5.2), imposing model constraints on the subdominant isotopes ^{238}U and ^{241}Pu, and fitting two separate models for ^{235}U and ^{239}Pu as scaled by the fission fraction, Daya Bay extracts the spectral components from each of the two main isotopes. In practice, this is done by minimizing χ^2_{DYB} defined in [83] as:

$$\chi^2_{DYB}(\eta^5,\eta^9) = 2\sum_{djk}\left(S^{DYB}_{djk} - M^{DYB}_{djk} + M^{DYB}_{djk}\ln\frac{M^{DYB}_{djk}}{S^{DYB}_{djk}}\right) + f(\epsilon,\Sigma), \tag{5.24}$$

where d is the index over 4 detectors, j is the index over 20 fission fraction data groups, and k is indexed over 26 prompt energy bins. η^5 are the fit parameters for the energy spectrum of ^{235}U, η^9 is correspondingly the parameters for ^{239}Pu (each with 26 independent fit parameters), ϵ is a set of nuisance parameters, and $f(\epsilon,\Sigma)$ is a term to constrain the nuisance parameters along with their correlations Σ. Finally, S^{DYB}_{djk} is the predicted spectrum such that

$$S^{DYB}_{djk} = \alpha_{jk}(\epsilon)s^5_k(\eta^5_k) + \beta_{jk}(\epsilon)s^9_k(\eta^9_k) + s^{238+241}_k(\epsilon) + c_k(\epsilon), \tag{5.25}$$

where $s^5_k(\eta^5_k)$ is the fitted spectrum for ^{235}U in Daya Bay's prompt energy space ($s^9_k(\eta^9_k)$ is similarly for ^{239}Pu), α_{jk} (β_{jk}) is the coefficient for s^5_k (s^9_k) accounting for absolute flux parameters, $s^{238+241}_k$ is the expected prompt energy spectrum for the subdominant isotope

fuel (as constrained by the Huber model), and c_k is the predicted spectrum from other backgrounds such as spent nuclear fuel. By minimizing χ^2_{DYB} over the fit parameter of η^5 and η^9, Daya Bay is able to generate its two deconvolved spectra.

Introduction of the PROSPECT data contributes an additional handle on the deconvolution via an additional data group with a ^{235}U fission fraction of 1.0. Using the same formalism as in Equations 5.24 and 5.25 an additional test statistic χ^2_{PRO} is defined:

$$\chi^2_{PRO} = 2\sum_i \left(S_i^{PRO} - M_i^{PRO} + M_i^{PRO} \ln \frac{M_i^{PRO}}{S_i^{PRO}} \right) + g(\delta, \Omega), \tag{5.26}$$

$$S_i^{PRO} = \gamma_i(\delta) \left(R^{map} \times s^5(\eta^5) \right)_i + d_i(\delta), \tag{5.27}$$

where M^{PRO} is the PROSPECT prompt energy measurement, $g(\delta, \Omega)$ is a term to constrain nuisance parameters δ with covariance Ω, $R^{map} \times s^5(\eta^5)$ is the mapping of the ^{235}U spectrum fitted in Daya Bay's prompt energy space into PROSPECT's prompt energy space (defined the same way as in Equation 5.16), d is the predicted contributions from other components in the PROSPECT measurement, γ is a scaling coefficient, and i is the index over 32 energy bins. Then by minimizing the sum of $\chi^2_{DYB} + \chi^2_{PRO}$ the Daya Bay LEU spectrum is deconvolved using the PROSPECT HEU measurement in Daya Bay's prompt energy space. The results from this deconvolution are shown in Figure 5.15 [84]. In this figure, the Huber model prediction for ^{235}U and ^{239}Pu are included and scaled to match the integrated rate of each spectrum for reference. The ratio of the unfolded data to the scaled models are given in the lower panel to highlight the differences in their spectral shapes.

The improvements from this joint deconvolution are shown in Figure 5.16 [84]. The relative changes in the central values of the two spectra are given in the top panel of the figure. The middle panel shows that while the relative uncertainty in the ^{239}Pu spectrum is largely unchanged, the relative uncertainty of ^{235}U improves by $\sim$0.5% across the entire

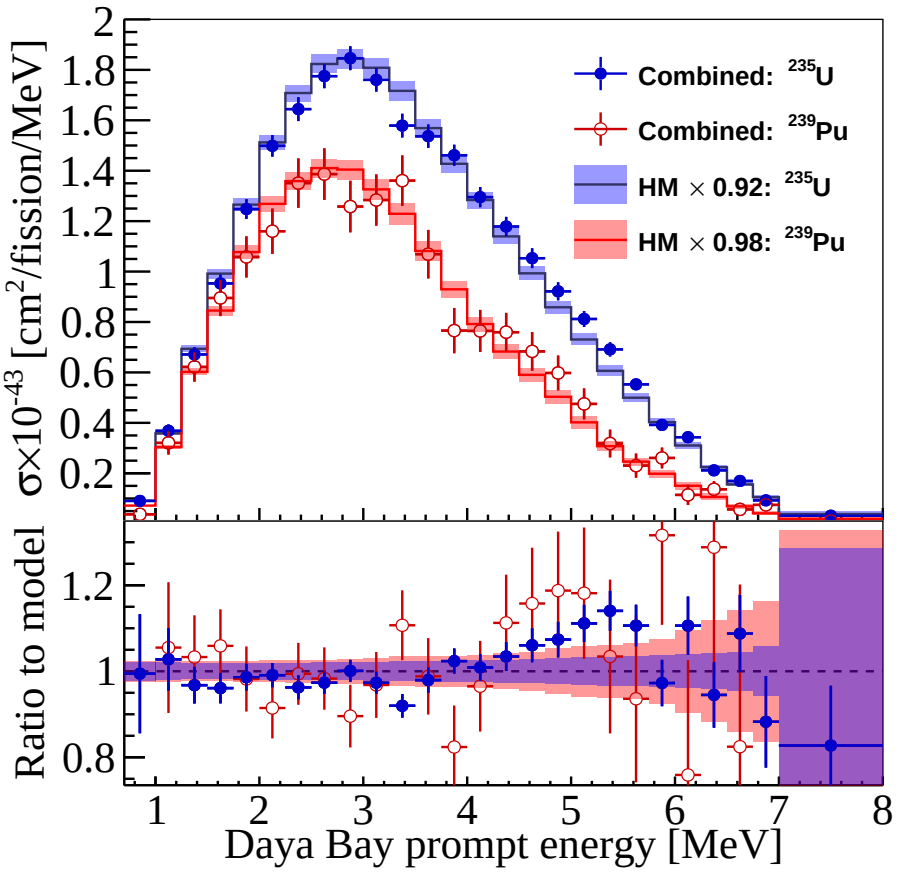

Figure 5.15: (top) Spectra for ^{235}U and ^{239}Pu in Daya Bay's prompt energy space as jointly deconvolved from the Daya Bay LEU measurement with the PROSPECT HEU measurement. The Huber-Mueller prediction for each spectrum is scaled to match the total rate, and included for reference. (bottom) Ratio of each measured spectrum to its scaled prediction, highlighting deviations in spectral shape [84].

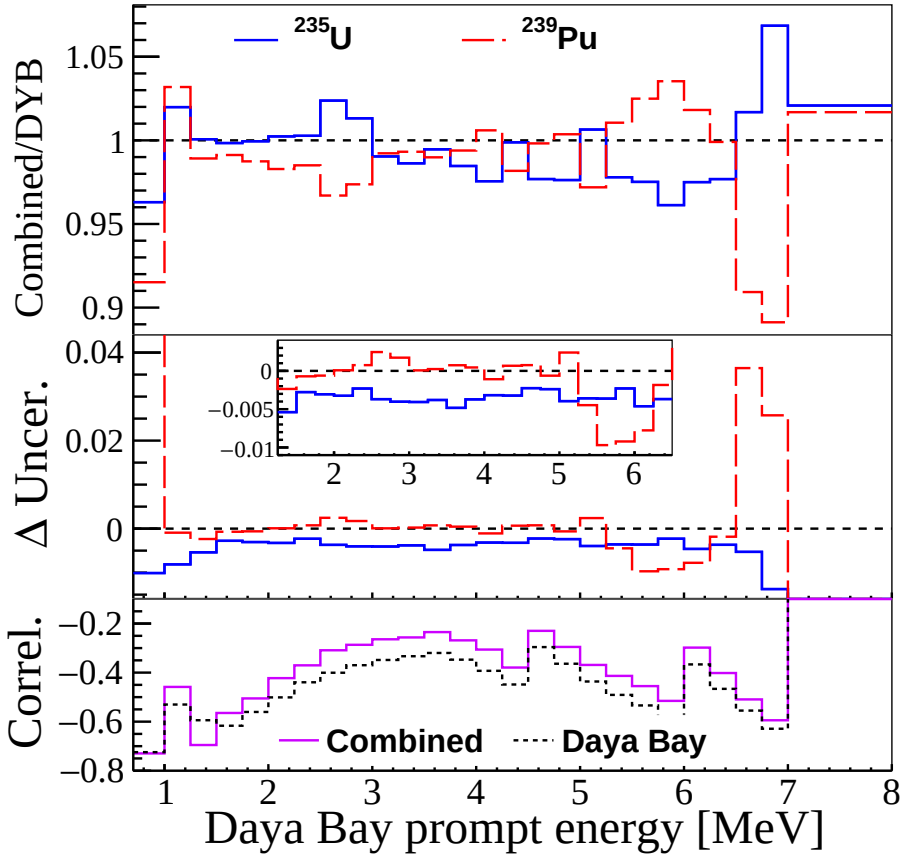

Figure 5.16: Changes and improvements of the jointly deconvolved LEU spectrum vs the Daya Bay only deconvolved spectrum as shown in Daya Bay's prompt energy space. (top) Relative changes in central values of the spectra. (middle) Relative change in uncertainty of each measurement, including an expanded view in the insert. While the ^{239}Pu uncertainty is largely the same, the ^{235}U uncertainty improves by approximately 0.5%. (bottom) Improved anti-correlation between the two deconvolved spectra, improving the degeneracy between the two by approximately 20% [84].

energy range (e.g. from 3.5% to 3.0% around 3 MeV). An expanded portion of the plot is given in the insert. The bottom panel of the plot shows the reduction in degeneracy between the two spectra, with an improvement of ∼20% across the whole energy range (e.g. anti-correlation changes from -0.5 to -0.4 near 2 MeV).

All three approaches to unfolding the ^{235}U spectrum yield consistent and comparable results, with the joint deconvolution method also improving on the relation between the ^{235}U and ^{239}Pu spectra. The first two methods, the joint fitting and joint unfolding, in-

dependently cross-check the joint deconvolution method. Possible future directions and ways to improve this measurement moving forward are included in Section 6.2

Chapter 6

Conclusions

6.1 Summary of Results

The PROSPECT experiment was designed to search for $\sim$1 eV^2 scale sterile neutrino signatures and to make a leading precision measurement of the ^{235}U $\bar{\nu}_e$ energy spectrum through inverse beta decay (IBD) interactions of reactor antineutrinos. After approximately two years of design and prototyping, and roughly another two years of construction and commissioning the PROSPECT detector, the 4 ton ^{6}Li-loaded liquid scintillator detector of 154 optically isolated segments achieved its technical design requirements in reconstructed position (5 cm in z-position) and energy resolution (4.5% at 1 MeV). Operating for 96 calendar days of reactor-on data taking at HFIR at an average baseline distance of 7.9 m, PROSPECT detected over 50,000 IBD events with a cosmogenic (accidental) signal-to-background ratio of 1.4 (1.8), an unprecedented achievement operating at the Earth's surface with little to no overburden.

To date, PROSPECT has produced leading results in both of its physics goals. PROSPECT has excluded a large portion of the allowed phase space for a sterile neutrino flavor, specifically excluding the historical RAA best fit point at 2.5σ. PROSPECT also made the first precision measurement of the $\bar{\nu}_e$ spectrum at a highly enriched uranium reactor. With re-

gards to the recent data-model discrepancies of reactor $\bar{\nu}_e$ spectra, PROSPECT disfavors the hypothesis that ^{235}U does not contribute to the excess around 6 MeV at 2.2σ as well as the hypothesis that ^{235}U is the sole contributor to the excess at 2.4σ.

The spectral measurement by PROSPECT at a highly enriched uranium reactor is found to be consistent with the ^{235}U component of the deconvolved low-enriched uranium spectrum measurement done by the Daya Bay experiment (p-value = 0.75). These complementary measurements are jointly combined using several independent analysis frameworks with consistent results obtained by each. The joint analysis leads to an improved deconvolution of the two main contributors to the full power reactor spectrum, namely an improvement in the relative shape uncertainty of the ^{235}U spectrum in antineutrino energy by approximately 0.5% (e.g. an improvement from 3.5% to 3.0% at 3.0 MeV). Additionally the degeneracy between ^{235}U and ^{239}Pu is reduced, improving by $\sim$20% (e.g. anti-correlation changes from -0.5 to -0.4 near 2 MeV).

6.2 Future Outlook

While the full data set of PROSPECT has been collected, there are a number of ways to possibly improve both the effective statistics and detector response for an updated analysis. The changes of the detector performance over time, due to scintillator optical property degradation and failing segments within the detector, complicated the analysis. To deal with the varying performance of the detector, the simplest approach of scaling to the worst case of the detector at the end of its data taking was adopted. In particular, this meant that the reconstructed energy needed to be artificially smeared to match the energy resolution at the end of data taking, and segments that failed at any point were treated as if they were never active. Besides the first order effects of decreased energy resolution and reduced active detector volume, these also impacted the selection efficiency and purity of the final IBD event selection, particularly increasing the correlated nH and nC* backgrounds more

thoroughly discussed in Section 4.1.

By subdividing the data taking period in to discreet time periods, these negative impacts can be mitigated. Instead of having to smear out the energy resolution or remove segments from the analysis based on the entire data set, these corrections can be done individually for each time period. Then, the measurement from each time period can be treated as a separate measurement and unfolded using a procedure similar to the joint unfolding approach outlined in Section 5.5. Instead of treating each time-period measurement as independent, however, correlated uncertainties between segments must be accounted for similarly to how correlations across baselines were handled in Section 4.2.2. Studies must still be done to optimize the physics reach of subdividing the dataset, but improvements in both detector response, overall statistics, and background rejection from implementing this technique could lead to non-trivial improvements in the PROSPECT ^{235}U spectrum analysis, particularly in the sensitivity to an excess around 5 MeV in reconstructed energy, the location of one of the major backgrounds. This different data set analysis, as well as future implementation of analysis using single-ended segments (segments that were turned off in the analysis due to exactly one PMT failing) and machine learning techniques to better select IBD events and reject backgrounds, should lead to an improved search for sterile neutrino signatures and more sensitive interpretations of the ^{235}U antineutrino spectrum. Future analysis establishing an absolute efficiency for PROSPECT will also allow a flux-based analysis of HFIR, which may aid future work as well.

Additionally, these improvements can aid in the deconvolution of an LEU spectrum with Daya Bay. This updated PROSPECT spectrum can be combined with an future Daya Bay spectrum with $\sim$1000 days more data. In conjunction with this analysis, a separate joint analysis between the STEREO experiment [87, 88] and PROSPECT has been done [89]. STEREO is another liquid scintillator experiment measuring the antineutrino spectrum from a highly enriched uranium reactor with comparable statistics. A potential 3-way analysis between the measurements of Daya Bay, PROSPECT, and STEREO could poten-

tially lead to an even greater improvement of the LEU spectrum deconvolution, especially if the final analysis frameworks and data sets of each experiment are used. In principle, such an analysis could improve the measurement of the ^{239}Pu spectrum and even begin to probe the spectra of the sub-dominant power reactor isotopes ^{238}U and ^{241}Pu.

6.3 Personal Contribution

What follows is a list of my major personal contributions to the PROSPECT experiment:

- During the early design stages of PROSPECT, I participated in the design, construction, operation, and analysis of two of the prototype detectors, namely PROSPECT-20, a 20 L detector designed to test various configurations and optical properties of a single, rectangular segment. I made the same contributions to PROSPECT-50, a 2-segment miniature version of the PROSPECT inner detector designed to test production components for the full detector as well as events correlated between segments in a PROSPECT-like detector.

- I helped to design and test the PROSPECT radioactive source calibration system, including stand alone prototypes and an implemented tubing system in the PROSPECT-50 prototype. I developed and managed the annealing process for the source tubes for use in the full PROSPECT detector, and helped coordinate dedicated source calibration runs during operation of the full detector.

- Throughout production of the internal components of the PROSPECT detector, I was a member of the main PMT module production team at Yale. I was a significant worker on all aspects of the production including cleaning, assembly, quality checks, and was responsible for training and managing visiting collaborators on the PMT production line. In addition to those general duties, I also developed and ran

the optical and electronic characterization of over 300 PMT modules for use in the PROSPECT experiment.

- I ran one of the two parallel working groups during the inner detector assembly. I managed the team doing the final cleaning and quality checks of all components that were given to the second team for assembly. This included managing all PMT modules, calibration system components, mechanical support structures, and reflector panels. I also managed various additional tasks associated with assembling the outer components of the PROSPECT detector. This included working with the aluminum secondary container, tensioning cables, detector cables, and various interfaces with the inner acrylic vessel.

- I spent time on site at Oak Ridge National Laboratory preparing for the arrival of the detector, and I designed and built the LiLS filling system on site. This included developing the system for and managing the team that mixed 28 barrels of LiLS in an ISO tank, transferred the LiLS into the detector volume, and sampled the LiLS for future quality assessments. I also managed the team that removed the LiLS from the detector during decommissioning of the detector.

- I took both general shifts and on-call DAQ expert shifts for detector monitoring purposes throughout data taking. Across the 18 months of total data taking, I took 11(6) general (DAQ-expert) shifts for a total of 624 (504) hours where I was responsible for detector operations.

- As part of the data analysis push with first physics data, I co-led the team optimizing IBD event selection. Together with the other co-lead, we developed the figures of merit to be used, processed all the data used in the analysis, and compiled the report on the finalized IBD event selection cuts. We then reran this analysis for the full data set with minor changes implemented to account for the effects from the modified

detector response.

- I developed the analysis framework for and led the PROSPECT effort on the joint spectral analysis with Daya Bay. As part of this, I worked with the lead analyzer from Daya Bay as well as presented on PROSPECT's behalf at several Daya Bay meetings. I co-led the paper writing effort on this analysis, and managed the paper response throughout the review process.

- The base analysis framework I wrote for the Daya Bay joint analysis was also used in the STEREO joint analysis. I helped develop parts of that analysis in a consulting role.

- I presented at seven conferences on behalf of the PROSPECT collaboration, covering topics such as prototype development, calibration performance, joint analysis work, and an overview of short baseline sterile neutrino searches.

Bibliography

[1] W. Pauli. Dear radioactive ladies and gentlemen. *Phys. Today*, 31N9:27, 1978. 1

[2] Charles Drummond Ellis and W. A Wooster. The average energy of disintegration of radium E. *Proc. Roy. Soc. Lond. A*, 117(776):109–123, 1927. doi: 10.1098/rspa.1927.0168. xi, 2

[3] J. Chadwick. Possible Existence of a Neutron. *Nature*, 129:312, 1932. doi: 10.1038/129312a0. 2

[4] E. Fermi. An attempt of a theory of beta radiation. 1. *Z. Phys.*, 88:161–177, 1934. doi: 10.1007/BF01351864. 2

[5] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. Detection of the free neutrino: A Confirmation. *Science*, 124:103–104, 1956. doi: 10.1126/science.124.3212.103. 2, 10

[6] G. Danby, J.M. Gaillard, Konstantin A. Goulianos, L.M. Lederman, Nari B. Mistry, M. Schwartz, and J. Steinberger. Observation of High-Energy Neutrino Reactions and the Existence of Two Kinds of Neutrinos. *Phys. Rev. Lett.*, 9:36–44, 1962. doi: 10.1103/PhysRevLett.9.36. 3

[7] J. Sielaff. Observation of tau neutrino charged current interactions. In *36th Rencontres de Moriond on Electroweak Interactions and Unified Theories*, 5 2001. 3

[8] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson. Experimental Test of Parity Conservation in β Decay. *Phys. Rev.*, 105:1413–1414, 1957. doi: 10.1103/PhysRev.105.1413. 4

[9] T. D. Lee and Chen-Ning Yang. Question of Parity Conservation in Weak Interactions. *Phys. Rev.*, 104:254–258, 1956. doi: 10.1103/PhysRev.104.254. 4

[10] T. D. Lee and Chen-Ning Yang. Parity Nonconservation and a Two Component Theory of the Neutrino. *Phys. Rev.*, 105:1671–1675, 1957. doi: 10.1103/PhysRev.105.1671. 4

[11] B.T. Cleveland, Timothy Daily, Jr. Davis, Raymond, James R. Distel, Kenneth Lande, C.K. Lee, Paul S. Wildenhain, and Jack Ullman. Measurement of the solar electron neutrino flux with the Homestake chlorine detector. *Astrophys. J.*, 496:505–526, 1998. doi: 10.1086/305343. 5

[12] B. Pontecorvo. Neutrino Experiments and the Problem of Conservation of Leptonic Charge. *Sov. Phys. JETP*, 26:984–988, 1968. 5

[13] Ziro Maki, Masami Nakagawa, and Shoichi Sakata. Remarks on the unified model of elementary particles. *Prog. Theor. Phys.*, 28:870–880, 1962. doi: 10.1143/PTP.28.870. 5, 7

[14] B. Pontecorvo. Inverse beta processes and nonconservation of lepton charge. *Zh. Eksp. Teor. Fiz.*, 34:247, 1957. 5, 7

[15] Samoil M. Bilenky and B. Pontecorvo. Lepton Mixing and Neutrino Oscillations. *Phys. Rept.*, 41:225–261, 1978. doi: 10.1016/0370-1573(78)90095-9. 5

[16] Q.R. Ahmad et al. Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 87:071301, 2001. doi: 10.1103/PhysRevLett.87.071301. 7

[17] Karsten Heeger. *Model-Independent Measurement of the Neutral-Current Interaction Rate of Solar 8B Neutrinos with Deuterium in the Sudbury Neutrino Observatory*. PhD thesis, Washington U., Seattle, 2002. 7

[18] K. Eguchi et al. First results from KamLAND: Evidence for reactor anti-neutrino disappearance. *Phys. Rev. Lett.*, 90:021802, 2003. doi: 10.1103/PhysRevLett.90.021802. 7, 10

[19] Y. Fukuda et al. Evidence for oscillation of atmospheric neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998. doi: 10.1103/PhysRevLett.81.1562. 7

[20] P. A. Zyla et al. Review of Particle Physics. *PTEP*, 2020(8):083C01, 2020. doi: 10.1093/ptep/ptaa104. 7

[21] K. Abe et al. Solar Neutrino Measurements in Super-Kamiokande-IV. *Phys. Rev. D*, 94(5):052010, 2016. doi: 10.1103/PhysRevD.94.052010. xxvii, 8

[22] A. Gando et al. Reactor On-Off Antineutrino Measurement with KamLAND. *Phys. Rev. D*, 88(3):033001, 2013. doi: 10.1103/PhysRevD.88.033001. xxvii, 8

[23] D. Adey et al. Measurement of the Electron Antineutrino Oscillation with 1958 Days of Operation at Daya Bay. *Phys. Rev. Lett.*, 121(24):241805, 2018. doi: 10.1103/PhysRevLett.121.241805. xxiii, xxvii, 8, 102, 103

[24] G. Bak et al. Measurement of Reactor Antineutrino Oscillation Amplitude and Frequency at RENO. *Phys. Rev. Lett.*, 121(20):201801, 2018. doi: 10.1103/PhysRevLett.121.201801.

[25] Y. Abe et al. Measurement of θ_{13} in Double Chooz using neutron captures on hydrogen with novel background rejection techniques. *JHEP*, 01:163, 2016. doi: 10.1007/JHEP01(2016)163.

[26] F. P. An et al. New measurement of θ_{13} via neutron capture on hydrogen at Daya Bay. *Phys. Rev. D*, 93(7):072011, 2016. doi: 10.1103/PhysRevD.93.072011. xxvii, 8

[27] A. Osipowicz et al. KATRIN: A Next generation tritium beta decay experiment with sub-eV sensitivity for the electron neutrino mass. Letter of intent. 9 2001. 8

[28] D. M. Asner et al. Single electron detection and spectroscopy via relativistic cyclotron radiation. *Phys. Rev. Lett.*, 114(16):162501, 2015. doi: 10.1103/PhysRevLett.114.162501. 8

[29] M. Aker et al. Improved Upper Limit on the Neutrino Mass from a Direct Kinematic Method by KATRIN. *Phys. Rev. Lett.*, 123(22):221802, 2019. doi: 10.1103/PhysRevLett.123.221802. 8

[30] Ali Ashtari Esfahani et al. Determining the neutrino mass with cyclotron radiation emission spectroscopy—Project 8. *J. Phys. G*, 44(5):054004, 2017. doi: 10.1088/1361-6471/aa5b4f. 8

[31] G. Anton et al. Search for Neutrinoless Double-β Decay with the Complete EXO-200 Dataset. *Phys. Rev. Lett.*, 123(16):161802, 2019. doi: 10.1103/PhysRevLett.123.161802. 9

[32] D. Q. Adams et al. Improved Limit on Neutrinoless Double-Beta Decay in ^{130}Te with CUORE. *Phys. Rev. Lett.*, 124(12):122501, 2020. doi: 10.1103/PhysRevLett.124.122501. 9

[33] R. N. Mohapatra and A. Y. Smirnov. Neutrino Mass and New Physics. *Ann. Rev. Nucl. Part. Sci.*, 56:569–628, 2006. doi: 10.1146/annurev.nucl.56.080805.140534. 9

[34] G. Mention, M. Fechner, Th. Lasserre, Th. A. Mueller, D. Lhuillier, M. Cribier, and A. Letourneau. The Reactor Antineutrino Anomaly. *Phys. Rev. D*, 83:073006, 2011. doi: 10.1103/PhysRevD.83.073006. xi, 10, 11, 81

[35] A. Aguilar-Arevalo et al. Evidence for neutrino oscillations from the observation of $\bar{\nu}_e$ appearance in a $\bar{\nu}_\mu$ beam. *Phys. Rev. D*, 64:112007, 2001. doi: 10.1103/PhysRevD.64.112007. 10

[36] A. A. Aguilar-Arevalo et al. A Search for Electron Neutrino Appearance at the $\Delta m^2 \sim 1eV^2$ Scale. *Phys. Rev. Lett.*, 98:231801, 2007. doi: 10.1103/PhysRevLett.98.231801. 10

[37] Carlo Giunti and Marco Laveder. Statistical Significance of the Gallium Anomaly. *Phys. Rev. C*, 83:065504, 2011. doi: 10.1103/PhysRevC.83.065504. 10

[38] S. Schael et al. Precision electroweak measurements on the Z resonance. *Phys. Rept.*, 427:257–454, 2006. doi: 10.1016/j.physrep.2005.12.006. 11

[39] Feng Peng An et al. Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay. *Phys. Rev. Lett.*, 116(6):061801, 2016. doi: 10.1103/PhysRevLett.116.061801. [Erratum: Phys.Rev.Lett. 118, 099902 (2017)]. 12

[40] Feng Peng An et al. Improved Measurement of the Reactor Antineutrino Flux and Spectrum at Daya Bay. *Chin. Phys. C*, 41(1):013002, 2017. doi: 10.1088/1674-1137/41/1/013002. xi, 12, 13, 99

[41] Y. Abe et al. Improved measurements of the neutrino mixing angle θ_{13} with the Double Chooz detector. *JHEP*, 10:086, 2014. doi: 10.1007/JHEP02(2015)074. [Erratum: JHEP 02, 074 (2015)]. xi, 12, 13, 94

[42] J. H. Choi et al. Observation of Energy and Baseline Dependent Reactor Antineutrino Disappearance in the RENO Experiment. *Phys. Rev. Lett.*, 116(21):211801, 2016. doi: 10.1103/PhysRevLett.116.211801. xi, 12, 13, 94

[43] D.A. Dwyer and T.J. Langford. Spectral Structure of Electron Antineutrinos from Nuclear Reactors. *Phys. Rev. Lett.*, 114(1):012502, 2015. doi: 10.1103/PhysRevLett.114.012502. 12

[44] Patrick Huber. On the determination of anti-neutrino spectra from nuclear reactors. *Phys. Rev. C*, 84:024617, 2011. doi: 10.1103/PhysRevC.85.029901. [Erratum: Phys.Rev.C 85, 029901 (2012)]. 14, 79, 95

[45] A. A. Hahn, K. Schreckenbach, G. Colvin, B. Krusche, W. Gelletly, and F. Von Feilitzsch. Anti-neutrino Spectra From ^{241}Pu and ^{239}Pu Thermal Neutron Fission Products. *Phys. Lett. B*, 218:365–368, 1989. doi: 10.1016/0370-2693(89)91598-0. 14

[46] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch. DETERMINATION OF THE ANTI-NEUTRINO SPECTRUM FROM U-235 THERMAL NEUTRON FISSION PRODUCTS UP TO 9.5-MEV. *Phys. Lett. B*, 160:325–330, 1985. doi: 10.1016/0370-2693(85)91337-1.

[47] K. Schreckenbach, G. Colvin, W. Gelletly, and F. Von Feilitzsch. Determination of the antineutrino spectrum from 235u thermal neutron fission products up to 9.5 mev. *Physics Letters B*, 160(4):325–330, 1985. ISSN 0370-2693. doi: https://doi.org/10.1016/0370-2693(85)91337-1. URL `https://www.sciencedirect.com/science/article/pii/0370269385913371`. 14

[48] Th. A. Mueller et al. Improved Predictions of Reactor Antineutrino Spectra. *Phys. Rev. C*, 83:054615, 2011. doi: 10.1103/PhysRevC.83.054615. 14, 94

[49] A. C. Hayes and Petr Vogel. Reactor Neutrino Spectra. *Ann. Rev. Nucl. Part. Sci.*, 66:219–244, 2016. doi: 10.1146/annurev-nucl-102115-044826. xi, 14, 15

[50] P. Vogel and John F. Beacom. Angular distribution of neutron inverse beta decay, anti-neutrino(e) + p —> e+ + n. *Phys. Rev. D*, 60:053003, 1999. doi: 10.1103/PhysRevD.60.053003. 14

[51] Carlo Bemporad, Giorgio Gratta, and Petr Vogel. Reactor Based Neutrino Oscillation Experiments. *Rev. Mod. Phys.*, 74:297, 2002. doi: 10.1103/RevModPhys.74.297. xi, 15, 16

[52] `https://prospect.yale.edu/`. 18

[53] `https://www.ornl.gov/`. 18, 21

[54] J. Ashenfelter et al. The PROSPECT Reactor Antineutrino Experiment. *Nucl. Instrum. Meth. A*, 922:287–309, 2019. doi: 10.1016/j.nima.2018.12.079. xiii, xiv, xvi, 18, 30, 31, 32, 33, 36, 37, 48, 49, 50, 51, 52

[55] M. Andriamirado et al. Improved short-baseline neutrino oscillation search and energy spectrum measurement with the PROSPECT experiment at HFIR. *Phys. Rev. D*, 103(3):032001, 2021. doi: 10.1103/PhysRevD.103.032001. xii, xv, xvi, xvii, xviii, xix, xx, xxi, xxii, xxiii, xxvii, xxviii, 18, 19, 21, 24, 43, 44, 45, 46, 48, 50, 53, 54, 55, 57, 58, 59, 62, 63, 64, 65, 66, 67, 70, 72, 73, 74, 75, 78, 79, 80, 81, 82, 84, 86, 87, 88, 89, 91, 95, 96, 97, 98, 99

[56] J. Ashenfelter et al. The PROSPECT Physics Program. *J. Phys. G*, 43(11):113001, 2016. doi: 10.1088/0954-3899/43/11/113001. 20

[57] `https://neutrons.ornl.gov/hfir`. 21

[58] A. B. Balantekin et al. Nonfuel antineutrino contributions in the ORNL High Flux Isotope Reactor (HFIR). *Phys. Rev. C*, 101(5):054605, 2020. doi: 10.1103/PhysRevC.101.054605. xii, 21, 22, 23, 94

[59] Christopher John Werner, Jeffrey S. Bull, C. J. Solomon, Forrest B. Brown, Gregg Walter McKinney, Michael Evan Rising, David A. Dixon, Roger Lee Martz, Henry G. Hughes, Lawrence James Cox, Anthony J. Zukaitis, J. C. Armstrong, Robert Arthur Forster, and Laura Casswell. Mcnp version 6.2 release notes. 2 2018. doi: 10.2172/1419730. URL `https://www.osti.gov/biblio/1419730`. 21, 94

[60] `https://eljentechnology.com/products/liquid-scintillators/ej-301-ej-309`. 23

[61] J. Ashenfelter et al. Lithium-loaded Liquid Scintillator Production for the PROSPECT experiment. *JINST*, 14(03):P03026, 2019. doi: 10.1088/1748-0221/14/03/P03026. xii, 23, 24, 25, 26, 27

[62] J. Ashenfelter et al. First search for short-baseline neutrino oscillations at HFIR with PROSPECT. *Phys. Rev. Lett.*, 121(25):251802, 2018. doi: 10.1103/PhysRevLett.121.251802. 23

[63] J. Ashenfelter et al. Performance of a segmented ^{6}Li-loaded liquid scintillator detector for the PROSPECT experiment. *JINST*, 13(06):P06023, 2018. doi: 10.1088/1748-0221/13/06/P06023. 24

[64] J. Ashenfelter et al. A Low Mass Optical Grid for the PROSPECT Reactor Antineutrino Detector. *JINST*, 14(04):P04014, 2019. doi: 10.1088/1748-0221/14/04/P04014. xiii, 27, 28, 29, 30, 31

[65] Michael F. Weber, Carl A. Stover, Larry R. Gilbert, Timothy J. Nevitt, and Andrew J. Ouderkirk. Giant birefringent optics in multilayer polymer mirrors. *Science*, 287 (5462):2451–2456, 2000. ISSN 0036-8075. doi: 10.1126/science.287.5462.2451. URL https://science.sciencemag.org/content/287/5462/2451. 29

[66] https://www.hitachi-metals.co.jp/e/products/elec/tel/pdf/hl-fm4-k.pdf. 32

[67] https://www.icotek.com/en-us. 39

[68] J. Ashenfelter et al. Background Radiation Measurements at High Power Research Reactors. *Nucl. Instrum. Meth. A*, 806:401–419, 2016. doi: 10.1016/j.nima.2015.10.023. 39, 74

[69] J. Ashenfelter et al. The radioactive source calibration system of the PROSPECT reactor antineutrino detector. *Nucl. Instrum. Meth. A*, 944:162465, 2019. doi: 10.1016/j.nima.2019.162465. xv, 41, 43

[70] https://www.caen.it/products/v1725/. 47

[71] http://www.wiener-d.com/sc/powered-crates/vme64-x/6u-vme64x-6023.html. 48

[72] https://7id.xray.aps.anl.gov/internal/hardware/nim/phillips_757.pdf. 48

[73] https://root.cern. 49

[74] https://www.sqlite.org/index.html. 49

[75] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli,

S. Giani, R. Giannitrapani, D. Gibin, J.J. Gmez Cadenas, I. Gonzlez, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F.W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampn, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M.G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. Di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J.P. Wellisch, T. Wenaus, D.C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche. Geant4?a simulation toolkit. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3): 250–303, 2003. ISSN 0168-9002. doi: https://doi.org/10.1016/S0168-9002(03)01368-8. URL `https://www.sciencedirect.com/science/article/pii/S0168900203013688`. 60

[76] J.B. BIRKS. Chapter 8 - organic liquid scintillators. In J.B. BIRKS, editor, *The Theory and Practice of Scintillation Counting*, International Series of Monographs in Electronics and Instrumentation, pages 269–320. Pergamon, 1964. ISBN 978-0-08-010472-0. doi: https://doi.org/10.1016/B978-0-08-010472-0.50013-6. URL `https://www.sciencedirect.com/science/article/pii/B9780080104720500136`. 60

[77] Gary J. Feldman and Robert D. Cousins. A Unified approach to the classical statistical analysis of small signals. *Phys. Rev. D*, 57:3873–3889, 1998. doi: 10.1103/PhysRevD.57.3873. 87

[78] X. Qian, A. Tan, J. J. Ling, Y. Nakajima, and C. Zhang. The Gaussian CL_s method for searches of new physics. *Nucl. Instrum. Meth. A*, 827:63–78, 2016. doi: 10.1016/j.nima.2016.04.089. 87, 93

[79] M. Andriamirado et al. Note on arXiv:2005.05301, 'Preparation of the Neutrino-4 experiment on search for sterile neutrino and the obtained results of measurements'. 6 2020. xxii, 88, 90

[80] Liang Zhan. Improved Measurement of Reactor Flux and Spectrum at Daya Bay. *J. Phys. Conf. Ser.*, 888(1):012132, 2017. doi: 10.1088/1742-6596/888/1/012132. 94

[81] Youngju Ko et al. Sterile neutrino search at the NEOS experiment. *PoS*, NuFact2017: 042, 2017. doi: 10.22323/1.295.0042.

[82] V. Zacek, G. Zacek, P. Vogel, and J. L. Vuilleumier. Evidence for a 5 MeV Spectral Deviation in the Goesgen Reactor Neutrino Oscillation Experiment. 7 2018. 94

[83] D. Adey et al. Extraction of the ^{235}U and ^{239}Pu Antineutrino Spectra at Daya Bay. *Phys. Rev. Lett.*, 123(11):111801, 2019. doi: 10.1103/PhysRevLett.123.111801. xxiii, 103, 105, 127

[84] F. P. An et al. Joint Determination of Reactor Antineutrino Spectra from ^{235}U and ^{239}Pu Fission by Daya Bay and PROSPECT. 6 2021. arXiv:2106.12251. xxiii, xxiv, xxvi, 104, 106, 117, 118, 128, 129, 130

[85] W. Tang, X. Li, X. Qian, H. Wei, and C. Zhang. Data Unfolding with Wiener-SVD Method. *JINST*, 12(10):P10002, 2017. doi: 10.1088/1748-0221/12/10/P10002. 105

[86] M. Hatlo, F. James, P. Mato, L. Moneta, M. Winkler, and A. Zsenei. Developments of mathematical software libraries for the LHC experiments. *IEEE Trans. Nucl. Sci.*, 52:2818–2822, 2005. doi: 10.1109/TNS.2005.860152. 120

[87] H. Almazán et al. First antineutrino energy spectrum from ^{235}U fissions with the STEREO detector at ILL. *J. Phys. G*, 48(7):075107, 2021. doi: 10.1088/1361-6471/abd37a. 134

[88] N. Allemandou et al. The STEREO Experiment. *JINST*, 13(07):P07009, 2018. doi: 10.1088/1748-0221/13/07/P07009. 134

[89] H. Almazán et al. Joint Measurement of the ^{235}U Antineutrino Spectrum by Prospect and Stereo. 7 2021. 2107.03371. 134