MEASUREMENT OF LOW-ENERGY NUCLEAR-RECOIL QUENCHING FACTORS IN CSI[NA] AND STATISTICAL ANALYSIS OF THE FIRST OBSERVATION OF COHERENT, ELASTIC NEUTRINO-NUCLEUS SCATTERING

Grayson Currie Rich

A dissertation submitted to the faculty at the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Physics.

Chapel Hill
2017

Approved by:

P.S. Barbeau

J. Engel

R. Henning

H.J. Karwowski

S. Washburn

ABSTRACT

Grayson Currie Rich: Measurement of low-energy nuclear-recoil quenching factors in CsI[Na] and statistical analysis of the first observation of coherent, elastic neutrino-nucleus scattering
(Under the direction of P.S. Barbeau)

The COHERENT Collaboration has produced the first-ever observation, with a significance of 6.7σ, of a process consistent with coherent, elastic neutrino-nucleus scattering (CEνNS) as first predicted and described by D.Z. Freedman in 1974 [111]. Physics of the CEνNS process are presented along with its relationship to future measurements in the arenas of nuclear physics, fundamental particle physics, and astroparticle physics, where the newly-observed interaction presents a viable tool for investigations into numerous outstanding questions about the nature of the universe. To enable the CEνNS observation with a 14.6-kg CsI[Na] detector, new measurements of the response of CsI[Na] to low-energy nuclear recoils, which is the only mechanism by which CEνNS is detectable, were carried out at Triangle Universities Nuclear Laboratory; these measurements are detailed and an effective nuclear-recoil quenching factor of $8.78 \pm 1.66\%$ is established for CsI[Na] in the recoil-energy range of 5–30 keV, based on new and literature data. Following separate analyses of the CEνNS-search data by groups at the University of Chicago and the Moscow Engineering and Physics Institute, information from simulations, calculations, and ancillary measurements were used to inform statistical analyses of the collected data. Based on input from the Chicago analysis, the number of CEνNS events expected from the Standard Model is 173 ± 48; interpretation as a simple counting experiment finds 136 ± 31 CEνNS counts in the data, while a two-dimensional, profile likelihood fit yields 134 ± 22 CEνNS counts. Details of the simulations, calculations, and supporting measurements are discussed, in addition to the statistical procedures. Finally, potential improvements to the CsI[Na]-based CEνNS measurement are presented along with future possibilities for COHERENT Collaboration, including new CEνNS detectors and measurement of the neutrino-induced neutron spallation process.

Rabbit faces the neutrino floor

ACKNOWLEDGEMENTS

Significant portions of this thesis were built upon the efforts of many members of The COHERENT Collaboration, but a few individuals are owed specific recognition:

Juan Collar The man, the myth, the footnotes, the unspeakable approaches to electrical grounding – Juan's non-material contributions are a bottomless bounty of science and art, kindness and fury. His material contributions, such as providing the CsI[Na] crystal used during the QF measurements presented here, are less poetic but of equal importance. My gratitude: ineffable.

Jason Newby Jason's tireless effort at ORNL and the SNS ensured that the COHERENT efforts could proceed so smoothly that some may fail to recognize the integral role he continues to play. Beyond the administrivia, Jason's technical savvy and the accompanying reputation, well earned, far proceeds him: I knew his name and his code years before I would join COHERENT. Beyond the know-how, he is quite simply disarmingly nice.

Yuri Efremenko Yuri's indefatiguable advancement of neutrino physics at the SNS laid the groundwork for COHERENT, and he has been a strong advocate for many efforts of importance to the work presented here.

Kate Scholberg Seemingly the nucleation site for COHERENT, Kate's leadership through the foggy first stages of collaboration have positioned COHERENT extremely well as it enters a kind of experimental adolescence.

Alexey Konovalov and Bjorn Scholz Sorcerers of scintillation, wizards of waveform analysis. Bjorn and Alexey made possible the observation of CEνNS and provide evidence that they still make some physicists like they used to.

Recognition and profound appreciation must be given to the technical staff of Triangle Universities Nuclear Laboratory: Jeff Addison, Bret Carlin, John Dunham, Patrick Mulkey, Richard O'Quinn, and Chris Westerfeldt. At various points, the sublime capabilities of the Duke University Instrument Shop and the helpfulness of its staff have been crucial. Bernie Jelnik, Phil Lewis, and Richard Nappi have made possible the realization of numerous experiments, (often) on short notice, and without too loudly complaining about my lack of planning.

The liquid scintillator detectors used in the CsI[Na] QF measurements presented in Chap. 4 were provided by Prof. John Mattingly of the Nuclear Engineering Department at North Carolina State University. Members of the Mattingly research group also provided invaluable assistance during the setup and execution of these measurements: Dr. Jonathan Mueller , Rob Weldon, and Mudit Mishra, were every bit as essential as the detectors they brought over.

In the course of analyzing neutron time-of-flight spectra and extracting a neutron beam energy distribution, this work used the Extreme Science and Engineering Discovery Environment (XSEDE), which is supported by National Science Foundation grant number ACI-1053575. Some additional TOF analysis and the development of the fitting software was carried out using the Duke Computing Cluster (DCC). Tom Milledge of Duke University provided extremely valuable guidance, support, and access to both of the aforementioned computing resources.

Jimmy Dorff and Barry Wilson have consistently and deftly aided through the years on a wide range of computing-related issues within the physics department at Duke; their expertise and willingness to offer it reduced the barriers that could be imposed by complicated computing needs. Barry and Jimmy also provided valuable connections to people beyond the physics department whose expertise proved valuable (indeed, it was through these two that I was connected to Tom and the DCC). Rounding out the team of people who let me `sudo` stupid things and don't yell when they have to clean up my mess are Alex Crowell and Bret Carlin.

The community of students, postdocs, and faculty at TUNL enhances the science and the experience of countless late nights. Prof. Diane Markoff is always a steady source of support and assistance. Ron Malone and Forrest Friesen have both contributed expertise, time, patience, deft tandem tuning, and earnest ears for bending. Dr. Jack Silano graciously donated his time to accelerator operation; Dr. Laurie Cumberbatch has an infinite reserve of kind words of support; and conversations with Aaron Swindell are always enriching. Dr. Zane Beckwith (of UNC, having avoided TUNL) is ultimately responsible for the creation of the UNC-Graduate-School-compliant LaTeX style file used to typeset this document.

The list of TUNL alumni who have indelibly brightened my research experience is long. Stephen Daigle, Johnny Cesaratto, and Paddy Finnerty made nuclear physics look much "cooler" than almost anyone and were amazingly welcoming. Prof. Richard Longland (alumni and back) is always a source of insightful conversation. Mitzi Boswell would always have an interesting idea or two lurking around and helped introduce me to a lot of the physics research lifestyle (she made logging in to check on simulation batch jobs look cool). Jeromy Tompkins and Charles Arnold provided a lot of the experience that would shape my drive to continue at TUNL and delve further into the experimental nuclear/particle physics world; I will never forgive them.

Prof. Hugon Karwowski has been a steadfast supporter, advocate, and voice of reason for many years.

Few are as skilled as Dr. K when it comes to diffusing the tension and frustration that accumulate during a difficult experiment. Prof. Tom Clegg has been a fixture of my time at UNC: he taught the first physics class I took, Fall semester of my freshman year. Embarrassingly, he learned my name, my face, and my inclination towards sleeping through class.

Tenzing Joshi, Kareem Kazkaz, and Samuele Sangiorgio of Lawrence Livermore National Lab (at one point or another) helped inform my thinking when it came to redirecting my research. My experience with them, beyond being enjoyable, made a lasting scientific impression and would inspire renewed passion and commitment to the pursuit of interesting physics.

My experience at LLNL would lead, ultimately, to my involvement with Phil Barbeau when he began his faculty position at Duke. The most meaningful things don't need embellishment: my time working with Phil has been transformative and significant. It can never be forgotten. The team, and my experience, was tremendously enriched by Long Li, Sam Hedges, Connor Awe, Katrina Miller, Ben Suh, as well as Justin Raybern, making cameo appearances.

We come now to the bit where there's no veneer of a professional relationship, and, again, I'm largely without words. These people have stoked the fires of passionate living: Marcus Medley, Scott Brennen, Kelly Clancy, Adrienne Heller, Zach Humphrey. Many people shaped my life for the better. I regret not being able to thank them all.

The longer I live apart from her, the more I realize that Nadja Vielot taught me much of what I know about being a socialized adult and what little I know about cooking. She also imprinted upon me an obsession with kitchen cleanliness, which makes me a better person.

The Pettigrew family (Charlene, Keith, and Stephen) have been welcoming and supportive through the latter stages of grad school, and they've been understanding of my increasing inattention to social interaction. Erica Pettigrew has been present through the vicissitudes of research. Her support, even when it seems futile, has been profoundly important.

There's a cat named Tokyo with whom I have shared emotionally significant times; he will appreciate this acknowledgment as much as he appreciates anything that isn't food. At the bottom of it all, my parents have always provided unquestioning, unconditional, unlimited encouragement and support. They have kept me going when nothing else can.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

LIST OF ABBREVIATIONS AND SYMBOLS

ADC	Analog-to-digital converter
BD	Backing detector
BPM	Beam pulse/pickoff monitor
CEνNS	Coherent, elastic neutrino-nucleus scattering
CFD	Constant-fraction discriminator
CMA	Conditional moving average
DAQ	Data acquisition (system)
FFT	Fast Fourier transform
FWHM	Full width at half maximum
GEANT4	Geometry and Tracking simulation toolkit, version 4
HOG	Hot off-gas
LIDAR	Light detection and ranging
(N)LL	(Negative) log likelihood
MCMC	Markov chain Monte Carlo
MCNP(X/6)	Monte Carlo n-particle transport code (version X/version 6)
(E)ML	(Extended) Maximum likelihood
NIN	Neutrino-induced neutron
OFHC	Oxygen-free, high-conductivity
ORNL	Oak Ridge National Laboratory
PDF	Probability-distribution function
(S,D/2,3,..)PE	(Single, double, triple, ...) photoelectron
PMT	Photomultiplier tube

POT	Protons on target
PS(D or P)	Pulse shape discrimination or parameter
PTFE	Polytetrafluoroethylene (teflon)
QF	Quenching factor
SASI	Standing accretion shock instability
(CC)SNe	(Core-collapse) supernovae
SNS	Spallation Neutron Source
SSA	Shielded Source Area
TOF	Time of flight
TUNL	Triangle Universities Nuclear Laboratory

CHAPTER 1: Coherent elastic neutrino-nucleus scattering (CEνNS)

Section 1.1: A simple prediction and "an act of hubris"

The introduction and formalization of the Standard Model of particle physics through the 1960's and 1970's laid the groundwork on which, after 50 years, we still base our understanding of the physics of fundamental particles. Much of the picture was drawn by the work of Glashow [121], Salam and Ward [218], and Weinberg [249] which unified the theories of the electromagnetic and weak-nuclear forces. Quite importantly, the new theory anticipated a neutral-current interaction mediated by a single boson whose mass was calculable, providing experimentally verifiable predictions such as the mass of the boson and, closely related, the interaction rates that might be expected by such a force [140, 250]. Though it would be some time before a particle with the predicted mass of the neutral-current-mediating boson would be detected by the UA1 [25] and UA2 [29] collaborations in 1983, a set of two papers from the Gargamelle experiment in 1973 presented evidence of neutral current interactions [128, 129].

The emergence of experimental support for the neutral current, still in advance of the UA1 and UA2 results, would secure Nobel prizes for Glashow, Weinberg, and Salam in 1979, but worked even more quickly to focus broad theoretical efforts on the framework presented by the (increasingly well-supported) Standard Model. Within 6 months of the publication of the Gargamelle observation, D.Z. Freedman reflected on the well-established process of coherent, elastic electron-nucleus scattering [134] and posited that the nascent weak neutral current should facilitate coherent scattering of neutrinos off of nuclei in rather analogous manner [111]. The nomenclature for the process predicted by Freedman is varied through literature, but is now referred to as coherent, elastic neutrino-nucleus scattering CEνNS[1].

The seeming simplicity of the CEνNS mechanism belies the significant swath of physics to which it is related. Continuing the analogy to electron-nucleus scattering, Freedman noted that CEνNS would be related to and could elucidate information about the weak-charge distribution of nuclei [111]. Freedman also observed that, as neutrinos are produced prolifically in many astrophysical settings, CEνNS is potentially of significance in both core-collapse supernovae, where it could provide an efficient coupling of neutrino energy

[1] CEνNS (pronounced "sevens") occasionally appears in literature as CENNS and the process has at other times been referred to simply as coherent neutrino scattering CNS. CEνNS provides additional specificity relative to CNS, distinguishing the process from higher-energy coherent neutrino-nucleus interactions, such as coherent pion production. The CEνNS acronym is slightly less ambiguous than CENNS as "NN scattering" could appear, to pious nuclear physicists, a reference to nucleon-nucleon scattering.

to the surrounding matter, and in neutron stars, where it may participate in nucleosynthesis and ejection of matter [111]. Moving beyond the initial prediction, CEνNS has long been a candidate detection method for "neutrino physics and astronomy" [94], offering a viable route for detection of neutrinos from supernovae which could not only reveal information about the explosion process but also about fundamental properties of the neutrino itself [136]. CEνNS can participate in open questions about the fundamental nature of the neutrino, including its electromagnetic properties and the related question of whether or not neutrinos are their own antiparticles, all issues linked to physics beyond the Standard Model [160]. Sensitivity to beyond-the-Standard-Model physics extends also to non-standard neutrino interactions (NSI) such as flavor-changing neutral currents [37], evidence for a new massive mediator Z' [37, 97, 171], as well as offering an avenue for a flavor-blind search for a theoretical sterile neutrino species [221]. Additionally, through this process neutrinos can become tools to study other physics; specifically, by examining the distribution of CEνNS event energies, one could gain a better understanding of neutron-rich matter [197] relevant to exotic, heavy nuclei and the structure of neutron stars [138]. As a final example of the importance of CEνNS, it should begin to be observable in future detectors dedicated to direct observation of WIMP dark matter, which will be an important test of those detectors as well as an irreducible background [82].

While introducing the idea of the process, Freedman noted presciently that seeking to experimentally observe CEνNS "may be an act of hubris" owing to numerous technical challenges; specifically, it was noted that the "most conspicuous and most difficult feature" of observing CEνNS is the fact that the single detectable result of such an interaction is a nucleus recoiling with very little kinetic energy [111]. The original paper also identifies interaction rates, perennially a challenge of neutrino measurements, and neutron backgrounds, as they produce nuclear recoils which are indistinguishable from those produced by CEνNS, as particular challenges [111]. Owing in no small part to the magnitude of these challenges, the CEνNS process eluded detection for nearly 45 years following its prediction despite the array of compelling physics insight it could provide.

Section 1.2: Cross section and N^2 dependence

CEνNS can be described as a neutrino scattering simultaneously with all of the nucleons in a nuclear target; the nucleons then recoil *in phase*, i.e. coherently, which gives rise to an increase in the scattering amplitude of the process, resulting in a predicted cross section proportional to the number of nucleons squared. Coherency, which provides an increased probability of scattering through purely quantum mechanical effects, restricts this process to generally lower-energy neutrinos. For a broad-scope view, consider the de Broglie wavelength λ_ν of an incident neutrino with energy E_ν: as E_ν increases, λ_ν will fall and eventually become

incomparably small with respect to the size of the scattering nucleus; for lower values of E_ν, however, the entire nucleus "fits" within λ_ν, and the neutrino can be imagined as probing all of the nucleons.

The process is more properly described in terms of the momentum transfer Q^2 from the neutrino to the scattering nucleus. Freedman proposed the differential cross section for a neutrino of energy E_ν scattering off a target nucleus of mass M should be given by [111]

$$\frac{d\sigma}{dQ^2} = \frac{G_F^2}{2\pi} \sin^4\theta_W \left(F\big(Q^2\big)\right)^2 \left[1 - Q^2 \frac{2ME_\nu + M^2}{4M^2E_\nu^2}\right],$$

where G_F is the Fermi constant $1.1663787(6) \times 10^{-5}$ GeV^{-2} [196], $\sin^2\theta_W$ is the Weinberg angle, and $F\big(Q^2\big)$ is the nuclear form factor which describes the spatial distribution of the neutrons and protons of the nucleus (see §1.4). For low energies, Freedman makes the simplification

$$\frac{d\sigma}{d\cos\theta} \approx \frac{G_F^2}{2\pi} \sin^4\theta_W A^2 E_\nu^2 \left(1 + \cos\theta\right),$$

for scattering of the neutrino into laboratory angle θ [111]. This result exposes the nascent knowledge of the weak force in 1974: the behavior $\frac{d\sigma}{d\cos\theta} \propto A^2$ suggests the cross section should scale like the square of the number of nucleons in the nucleus A (which enters via the form factor $F\big(Q^2\big)$). Though this conceptually fits with the simple picture used to describe the process, wherein a neutrino scatters coherently off of all nucleons, this uses an anachronistic picture of the Standard Model where neutrons and protons participate equally in the weak nuclear force.

The Standard Model provides that the weak charge of the proton is given by $Q_W^p = 1 - 4\sin^2\theta_W$. Existing measurements of $\sin^2\theta_W$ have thus far confirmed Standard Model predictions, placing $\sin^2\theta_W \approx 0.232$, resulting in a near-zero weak charge of the proton [196]. Consideration of the extremely small weak charge of the proton then makes an estimate of the CEνNS cross section proportional not to the number of *nucleons* A squared, but to the squared number of *neutrons* in the nuclear target, $\sigma \propto N^2$.

Incorporating knowledge gained in the intervening decades since the original calculation by Freedman, the trio of Barranco, Miranda, and Rashba derive the CEνNS cross section and include considerable detail, exposing sensitivity to a wide range of physics; the resulting differential cross section[2] is given in terms of the energy of the recoiling nucleus E_{rec}, the incident neutrino energy E_ν, and the mass of the target nucleus

[2]In the interest of explicitness and transparency, note that this expression for cross section (and others here) uses natural units. In this scheme, the mass terms are in units of keV and stewardship of the orphaned factor of $1/c^2$ is handled in the end, where the differential cross section appears in units of 1/keV^2 (/keV): to get the (differential) cross section in units apparently less "natural" but infinitely more practical, a multiplicative factor of $(\hbar c)^2 = 3.8938 \times 10^{-16}$ cm$^2 \cdot$ keV^2 is applied.

M by [37, 38]

$$\frac{d\sigma}{dT} = \frac{G_F^2 M}{2\pi}\left[(G_V + G_A)^2 + (G_V - G_A)^2\left(1 - \frac{T}{E_\nu}\right)^2 - \left(G_V^2 - G_A^2\right)\frac{MT}{E_\nu^2}\right]. \tag{1.1}$$

The terms G_V and G_A encode sums of interaction couplings over the quark content of nucleons. Maintaining generality, we can expand

$$G_V = \left[\left(g_V^p + 2\sum \epsilon_{\alpha\beta}^{uV} + \sum \epsilon_{\alpha\beta}^{dV}\right) Z + \left(g_V^n + \sum \epsilon_{\alpha\beta}^{uV} + 2\sum \epsilon_{\alpha\beta}^{dV}\right) N\right] F_{\text{nuclear}}^V\left(Q^2\right), \tag{1.2a}$$

$$G_A = \left[\left(g_A^p + 2\sum \epsilon_{\alpha\beta}^{uA} + \sum \epsilon_{\alpha\beta}^{dA}\right)\left(Z_\uparrow - Z_\downarrow\right) + \left(g_A^n + \sum \epsilon_{\alpha\beta}^{uA} + 2\sum \epsilon_{\alpha\beta}^{dA}\right)\left(N_\uparrow - N_\downarrow\right)\right] F_{\text{nuclear}}^A\left(Q^2\right), \tag{1.2b}$$

with: the number of spin up [down] neutrons (protons) in the target nucleus $N(Z)_{\uparrow[\downarrow]}$; the Standard Model neutral current vector (axial) coupling constants for protons $g_{V(A)}^p$; and the Standard Model neutral current vector (axial) coupling constants for neutrons $g_{V(A)}^n$. Another component in Eqs. (1.2) are the nuclear vector (axial) form factors $F_{\text{nuclear}}^{V(A)}(Q^2)$, which are functions of the momentum transfer Q; related to the spatial distribution of nuclear matter, the form factors also have the effect of "enforcing" the coherency requirement of CEνNS as the form factors will decrease at higher values of Q (see Sec. 5.7.1 for a more detailed discussion of form factors).

To work out the neutral-current vector and axial coupling constants in Eqs. (1.2), we will need to make use of the quark-level neutral-current parameters for ν-hadron scattering; from Ref. [44], these expressions are[3]

$$\xi_L(u) = \rho_{\nu N}^{NC}\left(\frac{1}{2} - \frac{2}{3}\hat{\kappa}_{\nu N}\hat{s}_Z^2\right) + \lambda_{uL}, \tag{1.3a}$$

$$\xi_L(d) = \rho_{\nu N}^{NC}\left(-\frac{1}{2} + \frac{1}{3}\hat{\kappa}_{\nu N}\hat{s}_Z^2\right) + \lambda_{dL}, \tag{1.3b}$$

$$\xi_R(u) = \rho_{\nu N}^{NC}\left(-\frac{2}{3}\hat{\kappa}_{\nu N}\hat{s}_Z^2\right) + \lambda_{uR}, \tag{1.3c}$$

$$\xi_R(d) = \rho_{\nu N}^{NC}\left(\frac{1}{3}\hat{\kappa}_{\nu N} + \hat{s}_Z^2\right) + \lambda_{dR}, \tag{1.3d}$$

with $\hat{s}_Z^2$ the Weinberg or weak mixing angle calculated in the "modified minimal subtraction scheme" $\overline{\text{MS}}$, $\hat{s}_Z^2 = \sin^2\theta_W = 0.23129(5)$ [196], and numerous radiative corrections whose values are collected in Tab. 1.1. Note that the values of $\hat{s}_Z^2$ and the radiative corrections are taken from different editions of the Particle Data

[3] In Table 10.3 of Ref. [44], where these expressions appear, ϵ is used rather than ξ. This replacement is made to avoid confusion with other uses of ϵ in the present context.

Correction	Value
$\rho_{\nu N}^{NC}$	1.0082
$\hat{\kappa}_{\nu N}$	0.9965
λ^{uL}	-0.0031
λ^{dL}	-0.0025
λ^{uR}	3.7×10^{-5}
λ^{dR}	$2\lambda^{uR}$

Table 1.1: Radiative corrections used in the cross section calculation for CEνNS. All values are taken from the Particle Data Group's *Review of Particle Physics* 2013 edition [44]. Subsequent editions of *Review* adopt a different format for expressing Standard Model constants, so the corrections are taken from the most recent edition which maintains consistency with References [37, 38, 160], which themselves use values given in contemporary *Reviews*.

Group *Review of Particle Physics*[4], the 2016 [196] and 2013 [44] editions respectively, in order to maintain consistency with the expressions in Barranco *et al.* [37, 38].

Making use of the parameters in Eqs. (1.3), proton and neutron couplings are then composed of appropriate sums over quark content and handedness: for vector coupling, g_V, the left- and right-handed contributions are added; for axial coupling, g_A, the right-handed terms are subtracted from the left in the case of neutrinos, while in the case of anti-neutrinos the left-handed terms are subtracted from the right. Working out these expressions, the vector couplings are given by

$$\begin{aligned} g_V^p &= 2\xi_L(u) + \xi_L(d) + 2\xi_R(u) + \xi_R(d) \\ &= \rho_{\nu N}^{NC}\left(\frac{1}{2} - 2\hat{\kappa}_{\nu N}\hat{s}_Z^2\right) + 2\lambda^{uL} + 2\lambda^{uR} + \lambda^{dL} + \lambda^{dR}, \end{aligned} \tag{1.4a}$$

$$g_V^n = -\frac{1}{2}\rho_{\nu N}^{NC} + \lambda^{uL} + \lambda^{uR} + 2\lambda^{dL} + 2\lambda^{dR}, \tag{1.4b}$$

and the axial couplings, for neutrinos, are

$$\begin{aligned} g_A^p &= 2\xi_L(u) + \xi_L(d) - [2\xi_R(u) + \xi_R(d)] \\ &= \frac{1}{2}\rho_{\nu N}^{NC} + 2\lambda_{uL} + \lambda_{dL} - 2\lambda_{uR} - \lambda_{dR}, \end{aligned} \tag{1.5a}$$

$$g_A^n = -\frac{1}{2}\rho_{\nu N}^{NC} + 2\lambda_{dL} + \lambda uL - 2\lambda_{dR} - \lambda_{uR}. \tag{1.5b}$$

[4]The value of $\hat{s}_Z^2$ in the 2013 *Review* is 0.2312 [44].

With the values from Tab. 1.1, Eqs. (1.4) and (1.5) evaluate to yield

$$\begin{aligned} g_V^p &= 0.03088, \\ g_V^n &= -0.5120, \\ g_A^p &= \pm 0.4952, \\ g_A^n &= \mp 0.5124, \end{aligned} \tag{1.6}$$

where the top (bottom) term corresponds to (anti)neutrino scattering.

The final constituents of Eqs. (1.2) are the ϵ parameters which correspond to non-standard interactions (NSI). Generically, these are written $\epsilon_{\alpha\beta}^{qP}$ where $q = u, d$ indicates the quark participating in the interaction, $P = L, R$ indicates the chirality of the involved neutrino, and α, β indicate incoming and outgoing flavor of the neutrino, taking on indices of e, μ, τ [37]. The notation used in Eqs. (1.2) is slightly simplified by grouping the vector NSI couplings $\epsilon_{\alpha\beta}^{qV} \equiv \epsilon_{\alpha\beta}^{qL} + \epsilon_{\alpha\beta}^{qR}$ and the axial couplings $\epsilon_{\alpha\beta}^{qA} \equiv \epsilon_{\alpha\beta}^{qL} - \epsilon_{\alpha\beta}^{qR}$. Non-standard interactions are broken into non-universal $\epsilon_{\alpha\alpha}^{qP}$ and flavor-changing $\epsilon_{\alpha\beta}^{qP}$, with $\alpha \neq \beta$, categories; accounting for both non-universal and flavor-changing varieties of NSI, the sums in Eqs. (1.2) run over $\alpha, \beta = e, \mu, \tau$.

With the careful treatment of G_V and G_A in Eqs. (1.2), the differential cross section for CEνNS admits sensitivity to a host of potential NSI and attention to the behavior as a function of proton and neutron number can allow design of experiments which are especially sensitive to certain kinds of NSI [221].

In certain circumstances, the cross section for CEνNS can be simplified slightly, allowing for calculations which are subject to fewer poorly constrained parameters [37, 221]. Coarsely, irrespective of the presence of NSI, the axial contribution to the cross section will generally be suppressed relative to the vector contribution by a factor of $\sim 1/N$ as the axial terms scale like the relative number of spin up/down protons or neutrons (i.e., axial terms are scaled by $0, \pm 1$) while the vector terms scale like the total number of protons or neutrons. Thus, an approximation of the Standard Model CEνNS rate, especially good for heavy nuclei, can be made by neglecting axial contributions. In the case of even-even nuclei, the vector-only approximation yields an exact value: the axial contributions to Eq. (1.1) vanish in the situation where the number of spin-up and -down neutrons (and protons) are equal. In the case of a spin-zero nucleus and neglecting any radiative corrections, the differential cross section given in Eq. (1.1) reduces to [19]

$$\frac{d\sigma}{dT} = \frac{G_F^2}{2\pi} M \left[2 - \frac{2T}{E_\nu} + \left(\frac{T}{E_\nu} \right)^2 - \frac{MT}{E_\nu^2} \right] \frac{Q_W^2}{4} \left(F\big(Q^2\big) \right)^2 . \tag{1.7}$$

Section 1.3: CEνNS and core-collapse supernovae

1.3.1: Participation in explosion process

Supernovae have been of interest to observers of the night sky for very nearly two millennia, with records of Chinese astronomers of the second century A.D. noting the appearance and fading of transient "guest stars" [49]. Two primary categories of explosion-type exists [146]: thermonuclear runaway and core collapse[5]. Even with a restricted focus to core-collapse scenarios, many sophisticated and purpose-built models do not yield successful explosions despite over 50 years of study on the subject [64].

During stellar core collapse, one of the general features thought to be present is an outward-moving shockwave that travels through the accreted material, dissociating nuclei and losing energy along the way [49]. This shockwave could, reductively, be considered *the explosion* of a dying star, with the preceding gravitational contraction a catalyst. The simple picture where a continuous shockwave blows off much of the stellar matter was problematic: when calculations were performed, the shockwave stalled in many cases and led to failed explosions [49, 63]. A possible solution to the stalled-shock problem is suggested by the tremendous amount of energy that is liberated during collapse in the form of neutrinos; a typical outward neutrino energy loss rate is $\sim 3\times10^{52}$ erg/s [49, 50]. Prior to the recognition of the CEνNS process, neutrinos were already being identified as a candidate for ensuring successful explosions by reinvigorating the stalled shock and driving the delayed shock mechanism [50, 72], and CEνNS provides an efficient avenue by which the energy contained in the neutrino flux could be redeposited in the matter surrounding the collapsed core; indeed, in the same paper in which the CEνNS process was first theoretically presented, its potential role in this and other astrophysical processes was identified as well [111]. Work in this vein continued [112, 226] and the "delayed shock" mechanism has remained a popular possible explanation over many successive supernovae literature reviews [49, 146, 256].

The importance of the role that CEνNS may play in supernovae may be influenced by ion-ion correlations in the core of a collapsing star, where densities exceed normal nuclear densities, $\rho \gtrsim 3\times10^{14}$ g/cm^{-3} [49]. As densities reach extremely high values, the scattering neutrino effectively begins to probe neighboring nuclei, and this results in a reduction of the cross section. Ion-ion correlations and their mitigation of the effect of CEνNS are mentioned as early as 1982 [57], but are omitted from the comprehensive review of supernova mechanisms by Bethe in 1990 [49]. Horowitz [135] explores the screening effect of ion correlations and its

[5]These two explosive drivers do not map cleanly between the "types" of supernovae, e.g. Type I or Type II: thermonuclear runaway is associated with Type Ia and core collapse with Type II and Type Ib,c. The types are a taxonomic code presented initially by Zwicky and determined by more observational elements, such as chemical composition or light yield; for instance, Type I supernovae are those without evidence of the presence of hydrogen, while Type II *do* show evidence of hydrogen. The types are not strictly observed by nature, however, with some SNe transitioning between types, such as SN1993j, which seemed to evolve from Type II to Type Ib [237].

impact in neutrino trapping during core collapse. Subsequent refinements to the study of this effect have showcased the level of detail which must be considered when calculating effects in these environments [143, 177].

Though the discussion here has focused on the role CEνNS could play in the mechanism of supernovae, the complexity of these astrophysical occurrences cannot be overstated and increasingly detailed simulations of the magnetohydrodynamics (MHD) of core collapse have presented a possible "competing" explanation for successful explosions in the form of the standing accretion shock instability (SASI) [55]. Whatever role CEνNS may have, the correlations that exist between different weak-nuclear effects in the CCSNe reaction network have been found, in simulation, to give rise to non-linearities in the response to variations in reaction rates [168]. This consequence is closely related to *Mazurek's law* which pertains to the strong feedback that exists in CCSNe between equations of state (EoS), MHD, and weak-nuclear effects [146, 163]. Even if the impact of CEνNS were included[6] in the study of Ref. [168], the non-linearities and correlations they observed in dynamics response from variation of different weak-nuclear rates highlights the difficulty in interpreting the results of computational studies of supernovae and inferring the significance of any particular reaction on the explosion mechanism. Numerical studies intended to "focus" on particular aspects of explosions, such as the work initially predicting SASI [55], have tended to make simplifications of other processes [64, 146]; similarly, Ref. [168] explored the effects of varied weak interaction rates and neutrino opacities but did not also include variation of MHD or EoS effects. In a world of finite computational resources, it is unfortunate that the nature of CCSNe may hide itself from view of even the most-slightly simplified model.

Neutrino-induced heating and/or convection have consistently been regarded as significant participants in the explosion mechanisms of core-collapse supernovae [49, 50, 72, 146]. Understanding of these extremely energetic environments is complicated on many levels, all of which are compounded by tight correlations between microphysics effects [146]. While the effect of CEνNS may be mitigated by ion-ion correlations in the dense environments of collapsed stars [135, 143, 177], a confirmation of the existence of the CEνNS process would nonetheless retire some uncertainty in the modeling of supernovae, though the quantitative impact is unlikely to be significant or easily interpretable.

1.3.2: Observation of supernova neutrinos

In addition to possible participation in the supernova explosion process, CEνNS and its relatively large cross section provide a mechanism for observation of supernova neutrinos [136]. Core-collapse supernovae are tremendously large many-body systems and they are driven predominantly by the weak force; this

[6] The authors of Ref. [168] do not vary the baseline CEνNS rate or include any effect of ion-ion correlation, adopting the community-canonical rate recommended by Ref. [60].

environment presents a unique laboratory for the study of fundamental physics topics including, but not limited to, neutrino flavor oscillations, the neutrino mass hierarchy, and neutrino-neutrino interactions [32]. Despite tight correlations between many uncertainties that go into models of CCSNe dynamics (see preceding section), some models of observables are somewhat decoupled from certain complexities [32], with the net result indicating that supernova neutrinos are viable lenses through which the CCSNe laboratory can be viewed.

During gravitational collapse, the extraordinary density and pressure leads to electron capture reactions which produce a large number of ν_e in the core [168]. As densities peak in the core, it becomes opaque to neutrinos, typically restricting the energy-outflow in ν_e to $\sim 10^{51}$ erg during this phase of the explosion [65]. At the point where nuclear densities are reached, the equation of state in the core stiffens, and a rebounding shockwave travels outward [168]. The number of neutrinos produced is tremendous: $\sim 10^{58}$ in total, carrying approximately 10^{53} erg. Remarkably, during certain stages of the collapse, neutrino-neutrino scattering processes are considered to be dominant and collective neutrino oscillation can occur [32]; these effects, combined with subsequent transport effects, can leave observable finger prints on properties of the released neutrinos. In a core collapse scenario, the flavor of neutrinos plays a significant role due to the electron-rich environment and the concentration of dissociated neutrons and protons: ν_e and $\overline{\nu}_e$ interact more readily with this matter via charged-current reactions than do other species of neutrinos. As a consequence of the charged-current reactions that $\nu_e/\overline{\nu}_e$ undergo during the supernova, the energy spectra of the neutrinos as they emerge and decouple energetically from the stellar environment are expected to show considerable flavor dependence, with electron flavors having a softer spectrum than ν_x (collectively, ν_μ, $\overline{\nu}_\mu$, ν_τ, $\overline{\nu}_\tau$) [65]. Propagation of the neutrinos through the very-dense CCSNe environment can lead to flavor oscillations subject to the MSW effect of matter-enhanced oscillation.

The tremendous number of outgoing neutrinos from a supernova carry information which could improve our understanding of both supernovae and neutrinos, themselves. Measurement of the timing distribution of SNe neutrinos could even provide evidence for the presence of hydrodynamic instabilities like the standing accretion shock instability (SASI), though the necessary count rates are likely to occur only in "megaton-class instruments" [175, 238]. Spectral characteristics can provide information on some of the more profound questions about the nature of neutrinos: neutrino oscillations [174] and the neutrino mass hierarchy [83, 95]. Ideas have been proposed for flavor-dependent neutrino thermometers relying on CC reactions [100], but CEνNS could provide an alternative and complimentary means of determining energy distributions [136]. As a NC process, CEνNS cannot distinguish between ν_e and ν_x interactions, however, when coupled with information from other supernova neutrino observatories that can provide constraints on spectral information for ν_e and $\overline{\nu}_e$, CEνNS is a very valuable tool by which ν_x energy spectra may be measured [136]. Using

CEνNS, it is expected that a 200-ton, liquid-neon detector system could enable extraction of $\nu_x \rightarrow \nu_e$ oscillation information by way of discerning between coarse models of the effective temperature for ν_x (T_{ν_x}) from a CCSNe at a distance of 10 kpc from Earth [136].

Section 1.4: Probe of nuclear structure

Despite being one of the most fundamental properties of a nucleus, the actual physical size of a nucleus remains a difficult property to fully constrain via measurement or theory. Owing to the non-zero electrical charge of the proton, the distribution of protons in nuclei is a relatively accessible quantity and has been measured with precision via techniques such as elastic electron scattering [87]. Neutron distribution measurements have been carried out using, for instance, pion scattering [119], but these approaches rely on models to extract results [197]. Other attempts have been carried out by taking advantage of parity-violating effects in electron scattering, wherein the asymmetry between scattering of left- and right-helicity electrons is used to probe the weak-charge distribution of a nucleus which is closely linked to the neutron distribution due to the small weak charge of the proton [139]. A measurement of the neutron radius of ^{208}Pb by the PREX collaboration [4] is an example of successful use of PVES, with a similar measurement planned for ^{48}Ca as the CREX experiment [137]. PVES presents a "cleaner" approach for making these measurements [197], but the number of data points remains low and these experiments are sizable undertakings.

There is considerable motivation for performing measurements of neutron distributions, as they provide input for nuclear structure models [197] and associated observables such as the neutron skin [209]. Improved understanding of the structure of neutron-rich nuclei can ultimately refine the equation of state for neutron matter, helping to constrain models for neutron stars [18, 138, 197]. Neutron stars are very energetic astrophysical objects which are not yet fully understood [164], but may have broad significance including connection to gravitational wave production [243]. Following core collapse, neutron stars represent one potential fate of exhausted stellar objects, and an improved understanding of the equation of state for neutron-rich matter could further our understanding of the inner workings of CCSNe and whether a star will leave behind a neutron star or a black hole [256].

Neutron distributions are also of considerable importance in precise tests of the Standard Model at low energies using atomic parity non-conservation [139]. Atomic parity violation (APV) experiments allow for very low-energy tests of the Standard Model and are sensitive to physics beyond the Standard Model in the form of additional neutral bosons [90] and possible leptoquarks [88]. Extremely precise measurements of APV effects have been carried out in ^{133}Cs [254] resulting in precise tests of SM predictions, but interpretation of increasingly accurate results from APV measurements will require improved knowledge of

neutron distributions in the experimental nuclei, as the atomic effects are ultimately sensitive to the overlap of electrons and neutrons [201–203].

Though Freedman observed that CEνNS could elucidate structural information in a way similar to electron scattering [111], Amanik and McLaughlin revitalized the discussion and suggested the use of CEνNS as an avenue by which neutron form factors, directly related to the neutron spatial distributions, could be measured [19]. The form factor $F\left(Q^2\right)$ is a function of the momentum transfer and is the Fourier transform of the nuclear matter distribution.Returning to Sec. 1.2, one can see that it appears in the expression(s) for the differential cross section, in terms of recoil energy, of the CEνNS process (i.e., Eqs. (1.1), (1.7)). Ref. [197] begins from (1.7) and expands the form factor in a way which makes the neutron and proton contributions separable; the proton and neutron form factors are then expanded in a Taylor series, ultimately leading to an expression for the neutron form factor in terms of the even moments of the neutron density distribution. Though spherical nuclei are the quantitative focus of Patton *et al.*, deformations can be accounted for in the expansions though their significance will generally be suppressed [197].

The authors of Ref. [197] consider detectors with targets of ^{40}Ar, ^{76}Ge, and ^{136}Xe, and it is predicted that the neutron radius could be measured to within "a few percent" using this technique. However, since the form factors appear in the differential cross section, their measurement relies on spectral information from the nuclear recoils induced by CEνNS interactions [19], and the "few percent" estimate assumes a 1% uncorrelated error on the detection efficiency for the CEνNS detector in use [197]. This assumption is significant and places a form factor measurement out of reach of 1$^{\text{st}}$-generation CEνNS experiments. Spectral performance is crucial to other physics results that could be extracted from CEνNS, so progress in this area is expected.

Section 1.5: Tests of the Standard Model and beyond

1.5.1: Measurement of $\sin^2\theta_W$ and sensitivity to "dark" Z bosons

The expression for the CEνNS differential cross section (1.1) shows dependence on the ubiquitous weak-mixing angle (or Weinberg angle) $\sin^2\theta_W$. A notable property of $\sin^2\theta_W$ is that it "runs": the observed value of $\sin^2\theta_W$ is dependent on the momentum transfer Q at which the observation is made, and deviations from the Standard Model predictions of $\sin^2\theta_W$ at any Q would indicate new physics. Several measurements exist at different Q and are shown in Fig. 1.1 [196]. Though some tension with the Standard Model predictions exist, no statistically significant deviation has yet been observed.

Utilizing a stopped-pion neutrino source (§2.5), CEνNS is a vehicle for a measurement of these parameters at a momentum transfer of $Q \sim 40$ MeV/c [14], comparable to but lower than the value of Q at which

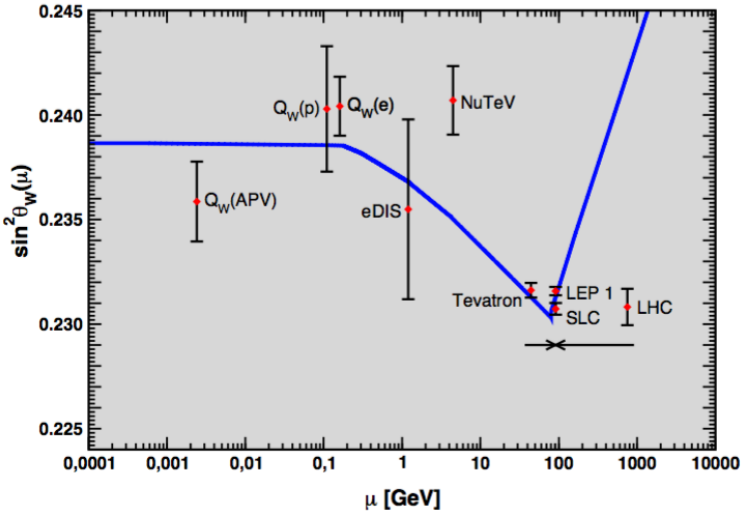

Figure 1.1: Existing measurements of $\sin^2\theta_W$ in the "modified minimal subtraction scheme" $\overline{\text{MS}}$ shown as a function of momentum transfer Q, stylized here as μ. Deviations from the Standard Model prediction, shown as a blue line, would indicated the presence of new physics. Using neutrinos from a stopped-pion source, with an energy distribution described in Section 2.5, CEνNS could provide a measurement of $\sin^2\theta_W$ at a momentum-transfer value of approximately 40 MeV/c, slightly below the region probed by proton and electron scattering experiments shown as $Q_W(p)$ and $Q_W(e)$, respectively. Measurements at the smallest value of Q come from cesium atomic parity violation experiments. Tevatron and LHC data points are shifted horizontally for clarity, but are taken at the same momentum transfer value as the LEP1 and SLC results. From Reference [196].

proton and electron scattering experiments make these measurements, see Fig. 1.1. A measurement of $\sin^2\theta_W$ produced by CEνNS would be subject to flux uncertainties and other systematics (e.g., detector response at low nuclear-recoil energies) at the level of $\sim 10\%$, likely rendering such a measurement non-competitive with the $\mathcal{O}(1\%)$ measurements produced by APV or PVES. Despite the comparatively large uncertainty, the distinct momentum transfer at which CEνNS could probe $\sin^2\theta_W$ would be attractive as the radiative corrections to the Standard Model are also Q dependent [139]. Additional value from a CEνNS-based measurement of $\sin^2\theta_W$ would come in the form of a unique set of systematic uncertainties, as the experimental probe is unique from those employed in other measurements.

Relationship to searches for dark Zs

Measurements of the muon anomalous magnetic moment a_μ have shown disagreement with the Standard Model prediction at levels $> 3\sigma$ (see Ref. [185] and references therein). Though not quite rising to the accepted 5σ level of "discovery", the significance of the observation has led to proposals for theoretical explanations of the measured values. One such proposal also presents a possible explanation for the phenomenon of dark matter (§1.6.1), and invokes the existence of as-yet unobserved low-mass gauge bosons called "dark Z" (Z_d) [85].

A manifestation of these dark Zs could be found in the running of the weak mixing angle: at low momentum transfers, $Q \lesssim m_{Z_d}$, the effective value of $\sin^2\theta_W$ is modified, which means that measurements of $\sin^2\theta_W$ at low Q^2 are sensitive to the existence of Z_d [86]. With its relatively low Q^2, CEνNS provides a probe sensitive to a larger range of Z_d masses than PVES measurements (see preceding discussion); however, as in the case of providing a measurement of $\sin^2\theta_W$, CEνNS tests of Z_d theories will be non-competitive with PVES or APV experiments for some time.

Spectral distortions from new mediators

Much of the earlier theoretical work surrounding new mediators, i.e. Z_d, focused on the gross scaling effect such a mediator might have on the CEνNS cross section, enacted via modification of $\sin^2\theta_W$. Liao and Marfatia considered the possible spectral implications a new mediator may have: by exploiting any nuclear-recoil spectral information that can be extracted from a CEνNS measurement, they find that the capabilities of this approach are enhanced relative to what might be extracted from a more-straightforward $\sin^2\theta_W$ measurement [171]. The authors of Ref. [171] work out the spectral effects in the case of CsI, but in considering the possible existence of an otherwise-undetected exotic mediator Z' they do not include axial contributions to the Standard Model CEνNS differential cross section. Though axial contributions will be suppressed considerably relative to vector (§1.2), they will be nonzero in a rigorous calculation of the CEνNS cross section for both Cs and I: Cs and I both have an unpaired proton so that the quantity $Z_\uparrow - Z_\downarrow$ in Eq. (1.2b) will be nonzero and yield a net nonzero value of G_A, the axial coupling strength entering the calculation of the differential CEνNS cross section. Nonetheless, the work of Liao and Marfatia [171] underscores the virtue of exploring spectral impacts of any exotic physics search.

1.5.2: Electromagnetic properties of the neutrino

Papavassiliou *et al.* show in Ref. [194] that a measurement of the nuclear recoil spectrum from CEνNS, already a subject of interest in §1.4, can inform an understanding of the effective neutrino charge radius. By measuring the CEνNS process with different neutrino flavors and the same target nucleus, cancellation of certain effects can produce an improved measurement of the difference of the effective charge radii of the flavors [194]. Flavor-dependent charge radii could result in different indices of refraction for neutrinos in matter [229] which may have astrophysical implications related to oscillation and propagation through both normal stellar matter and the dense environment during core collapse.

The magnetic moment of the neutrino μ_ν has very interesting implications if a measurement of its value could be performed. Most notably, massive Majorana and Dirac neutrinos have distinct allowed values of

μ_ν, so a measurement could provide insight into the Majorana vs. Dirac nature of the neutrino [150]. The capability to discern between these possible neutrino classes would be a tremendous boon, as it presents an independent avenue for confirmation of any potential observation of the $0\nu\beta\beta$ process by, for instance, the MAJORANA [198] or EXO [27] collaborations.

It has been recognized that the CEνNS recoil spectrum demonstrates sensitivity to the value of μ_ν at very low recoil energies [98, 221]. Reference [98] performs a careful sensitivity study of a hypothetical CEνNS measurement and the values of μ_ν that it is able to probe with given threshold values and exposures: with a 1 eV threshold for a Si/Ge detector, with 10^4 kg $\cdot$ yr exposure, *excluding* the effects of systematic uncertainties, this hypothetical CEνNS-based measurement of μ_ν is sensitive down to $\sim 2\times10^{-12}$ μ_B, where $\mu_B = e\hbar/(2m_e)$ is the Bohr magneton.

Current best limits on the neutrino magnetic moment are derived from solar neutrino data from the Borexino collaboration, which suggest $\mu_\nu < 1.4\times10^{-13}\mu_B$ [36]. The TEXONO and GEMMA experiments, which both utilize elastic scattering of reactor antineutrinos on electrons, have produced limits of $7.4\times10^{-11}\mu_B$ [253] and $2.9\times10^{-11}\mu_B$ [41], respectively. Based on the study of Ref. [98], it is clear that currently imagined CEνNS experiments can not probe beyond the best astrophysical limits on μ_ν. However, with the compelling implications that come with an understanding of the neutrino's electromagnetic properties, robust confirmation of these limits is still valuable. Additionally, it should be noted that the flavor composition of the neutrinos produced at stopped pion neutrino sources could allow direct determination of limits on the magnetic moment of muon neutrinos and antineutrinos.

1.5.3: Search for sterile neturinos

The picture presented by the Standard Model of 3 neutrino flavors is broadly compatible with most available experimental evidence, but several anomalous results have complicated this understanding. Two of the anomalous results, LSND [8] and MiniBooNE [9], are associated with neutrino beam measurements which detected evidence for an excess number of $\overline{\nu}_\mu \rightarrow \overline{\nu}_e$ transitions by observing $\overline{\nu}_e$ appearance. By contrast, the gallium anomaly, suggested by the gallium-based SAGE [3] and GALLEX [21] experiments, indicates a deficiency in the number of $\overline{\nu}_e$ events relative to 3-flavor predictions using radioactive ^{51}Cr sources. Finally, the reactor anomaly is based on the collective input of neutrino detectors which measure $\overline{\nu}_e$ produced at nuclear reactors and suggest a deficit relative to expectations from modern spectral calculations (see Ref. [182] for an overview).

To explain the observed anomalies, the existence of an additional species of neutrino which does not interact via the weak force has been posited [122]; the additional neutrino(s) are referred to as "sterile". A

model with a single additional sterile neutrino can explain the anomalous LSND and MiniBooNE results, in isolation, if the added neutrino has a squared mass difference of ~ 1 eV^2, assuming mixing between active and sterile species is similarly described as active-active oscillations [122]. Similarly, the reactor and gallium anomalies are accounted for with the introduction of a single, different sterile species [122]. Present analyses which incorporate results from the the many experiments do not find a resolution to the open issues and demonstrate tension when all anomalous results are simultaneously considered [122, 157].

Several different experiments have been proposed [20, 97, 109] which use CEνNS to search for evidence of sterile neutrinos. A CEνNS based sterile search would generally fall into the category of a disappearance experiment: any evidence for active-to-sterile oscillation would be manifested as a deficit of CEνNS events relative to the number expected. All previous efforts to investigate sterile neutrinos have used charged-current interactions [20], which introduce additional sources of uncertainty into any measurement. As a neutral-current process, CEνNS could present a powerful search tool for theoretical sterile neutrinos as it is sensitive to all active neutrino flavors. This flavor-independent approach theoretically allows for a sterile-neutrino search which encompasses the entire active-to-sterile oscillation space. Experimental realities make a CEνNS-based sterile-neutrino search challenging, but conceivable: Dutta *et al.* anticipate that an effective exposure of 2 years using germanium detectors with a target mass of 100 kg at a research reactor could probe much of the model space that is covered by existing short-baseline oscillation observations [97].

In addition to dedicated sterile-search experiments based on the CEνNS mechanism, Kosmas *et al.* [161] explored the sensitivity of other CEνNS experiments, making use of two distinct sources of neutrinos, to the existence of a sterile neutrino. For CEνNS experiments located at both nuclear reactors and stopped-pion sources (see a discussion of neutrino sources in §2.1.2), Ref. [161] find that sensitive CEνNS experiments could provide additional constraints to the sterile neutrino mixing-parameter space in a theoretical framework where a single sterile neutrino is present. Due to the difference in flavor composition, the constraint is more powerful if separate CEνNS observations are made using both neutrino sources. A notable distinction between the sensitivity found in Ref. [161] and dedicated sterile experiments is that the experimental configurations considered by Kosmas *et al.* do not collect data at different standoff distances from the neutrino sources.

Section 1.6: Irreducible background for direct dark matter searches

1.6.1: A new type of matter to explain anomalous observations

In 1970, observation of the rotation of the Andromeda Nebula [216] confirmed a long-standing, bold proposal by Zwicky and overturned any belief that the scientific community fully, or even *nearly*, understood

the laws and knew the contents of the universe: the observed rotation was entirely incompatible with the gravitational profile predicted by the observed matter in the galaxy. A broadly accepted possible explanation for this behavior is the existence of a previously unknown form of matter which does not generate light nor interact directly via the electromagnetic force, earning it the name "dark matter", first coined by Zwicky in the 1930's to explain observations of the Coma cluster [259, 260]. The astroparticle and cosmological physics communities have devoted considerable effort towards developing an understanding of dark matter and directly detecting it.

While numerous theoretical candidate particles exist, weakly-interacting massive particles (WIMPs) have received much attention from experimentalists seeking to make a pioneering direct observation of dark matter interactions. WIMPs gained popularity in part due to the fact that, in certain theoretical frameworks, the observed abundance of dark matter requires particle properties and interaction cross sections which are comparable to those predicted by supersymmetric extensions of the standard model [47, 149]; this coincidence is referred to as the "WIMP miracle". In another fortuitous coincidence, this very-popular theory was compatible with a conceptually straight-forward experimental signature: the scattering of WIMP dark matter by a nucleus should, like CEνNS and electron-nucleus scattering, be subject to a coherency boost, and the recoiling nuclei should leave an observable signal.

1.6.2: Direct searches for WIMPs and the "neutrino floor"

The efforts toward direct detection of WIMP dark matter rely on the same mechanism of signal production as a CEνNS search: the low-energy recoils of nuclei in the detector volume resulting from coherent scattering of an incident particle[7][116]. With the underlying signal generation pathways being the same, and with the coherency requirement limiting recoil energies in both cases to the same regime, CEνNS interactions in WIMP detectors would be indistinguishable from dark matter interactions through pulse-shape discrimination, event topology, or other many other methods that reduce backgrounds in these detectors. Combined with the absence of timing structure to meaningfully separate the events, and in the absence of directional sensitivity [11], CEνNS interactions from solar or cosmological-background neutrinos represent an "irreducible" background in WIMP searches.

There are several neutrino sources expected to participate in potential CEνNS backgrounds for WIMP searches. The primary processes for the hydrogen burning that takes place in the sun are the proton-proton (pp) chains and the carbon-nitrogen-oxygen (CNO) cycles; certain reactions from both groups produce neutrinos which reach terrestrial experiments, including $p(p, e^+\nu)d$, electron capture on ^{7}Be, and β^+ decay

[7] The relationship between CEνNS and WIMP detectors is commented upon further in §2.2 and the mechanism by which signals are produced from these interactions, for a certain class of detector, is discussed in §3.1.

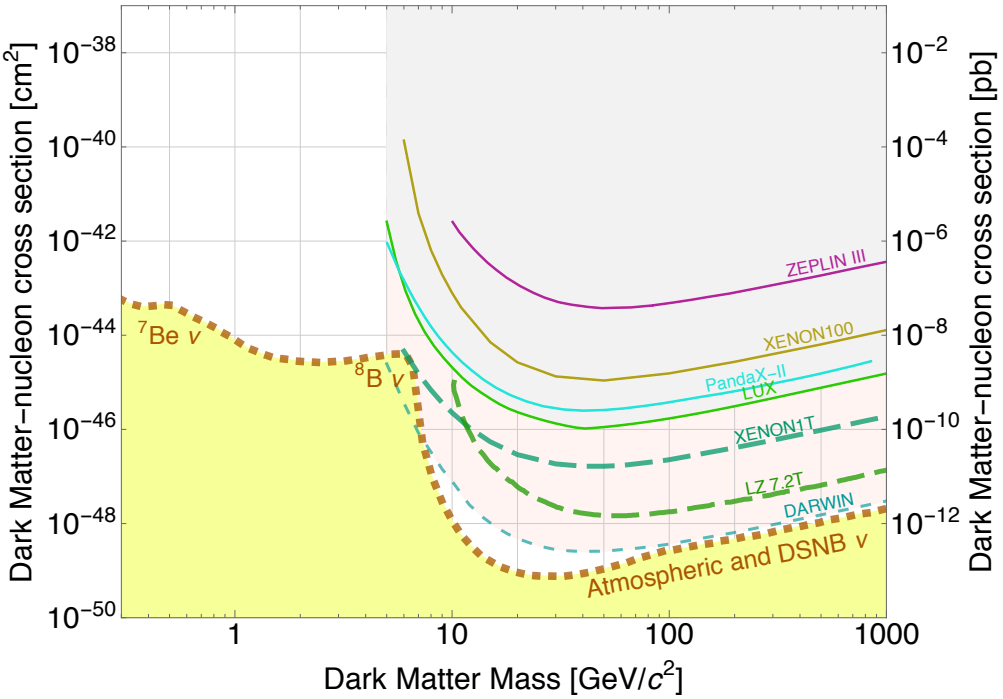

Figure 1.2: Plot of spin-independent WIMP search exclusion limits for existing (solid lines) and future (dashed lines) measurements. All shown experiments are based on xenon target nuclei. The so called "neutrino floor" can be seen towards lower cross sections. The location of the neutrino floor shown here assumes a Xe target. Figure created using Ref. [217].

of ^{8}B, ^{13}N, ^{15}O, and ^{17,18}F (see, e.g., Refs. [30, 141] for a more complete discussion). Neutrinos produced by cosmic rays in the atmosphere [115] provide another source, and the final expected source of CEνNS-background neutrinos come from the diffuse supernova neutrino background (DSNB), which is an isotropic distribution of the neutrinos produced by "all core-collapse supernovae in the causally reachable universe" [40]. The DSNB has yet to be experimentally observed, though limits on the flux have been placed by the Super-Kamiokande Collaboration [39].

DM direct detection experiments through the present generation have yet to achieve sufficient sensitivity to observe the expected CEνNS backgrounds. However, some of the planned next-generation experiments anticipate observation of CEνNS under the assumption of a standard-model cross section and understood astrophysical/cosmological neutrino fluxes and energies. Figure 1.2 shows the WIMP exclusion limits from xenon-based WIMP searches through the present generation and the anticipated sensitivities for several future experiments; also shown for a Xe target is the expected "neutrino floor", or the sensitivity surface at which CEνNS will become a measurable background.

Achieving sensitivity to CEνNS will mark a significant milestone in modern astroparticle physics, with a process that has eluded detection for several decades becoming a significant background due to progress in

technology and design. Though sensitivity to CEνNS will bring with it existential questions about further efforts towards larger, more sensitive WIMP detectors without directional-detection capability, CEνNS will serve as a uniquely valuable test for the performance of WIMP detectors, allowing for *in situ* confirmation of sensitivity estimates.

This serves to underscore the importance of an independent CEνNS observation and determination of the cross section. Absent confirmation of its existence and its cross section, failure to observe CEνNS in large WIMP detectors would be troubling and ambiguous. Such a scenario would leave only questions about the process itself and the workings of the detector and associated analysis. With considerable intellectual-, physical-, and financial-capital investment under way by the physics community in WIMP search efforts, independent measurement of the CEνNS process can serve as an invaluable cross check and confirmation of the performance of WIMP detectors at a very demanding level.

Section 1.7: Summary of this work

Motivated by the physics opportunities presented in the preceding sections, the work ultimately presented in this thesis was undertaken with the goals of enabling and participating in the first observation of the CEνNS process.

One of the enabling efforts was designing, constructing, and deploying a detector system intended to measure the neutrino-induced neutron (NIN) production process on ^{208}Pb. NINs and their importance, both as a background for a CEνNS measurement and as an important participant in numerous astrophysical (detection) scenarios, are discussed in various places through this work. A liquid-scintillator-based detection system was designed, simulated, built, characterized, and deployed to the Spallation Neutron Source (SNS) of Oak Ridge National Lab (ORNL). Section 7.8 covers much of the physics and presents a very brief summary of some of the work that was conducted towards a NIN measurement. The author played a primary role in all of these efforts.

The CEνNS observation discussed in this work relied only on determination that the NIN background was negligible (see §5.5). NIN data continues to be collected at the SNS with the neutrino cube detector(s) discussed in Sec. 7.8.2, and analysis will be an effort with which the author will be involved in the future, though likely not in a primary role.

Measurement of the quenching factor (QF) for nuclear recoils in CsI[Na] represents another, significant CEνNS-enabling effort. Chapter 4 describes an experiment carried out to perform such a measurement, and the synthesis of a representative value for the QF from the combination of new measurements and literature values. The author played a primary role in the planning, execution, and analysis of this measurement.

Additionally, the author was heavily involved in the execution and on-the-fly planning of a similar measurement, largely led by a team from the University of Chicago, whose result is also presented in Chap. 4. This additional result, included in Sec. 4.7, is the product of analysis with which the author was not involved.

Chapters 5 and 6 together describe the first observation of CEνNS, which was published as Ref. [17]. This result is the product of significant effort from numerous collaborators, but special recognition must be given to J.I. Collar of the University of Chicago. The author developed and carried out early simulations of the CEνNS experiments planned as a part of the COHERENT Collaboration; these simulations do not appear directly in this thesis, but confirmed independent findings, influenced general design decisions, and were included in a proposal submitted to the U.S. Department of Energy by the collaboration. The author was engaged with many discussions through the collection and analysis of the CsI[Na] CEνNS data, though the analysis credit generally belongs to B. Scholz and A. Konovalov. Simulations of the experimental geometries, discussed in §§5.6, 5.5, were troubleshot and run by the author, using MCNPX-PoliMi input files developed by collaborators.

Development and execution of the statistical analyses presented in Chap. 6 was the responsibility of the author, in close collaboration with P.S. Barbeau and J. Detwiler. This analysis incorporated CEνNS recoil distributions produced by P.S. Barbeau and subsequently investigated independently by the author, as well as analysis efficiency curves developed by the aforementioned analyzers.

CHAPTER 2: The COHERENT effort to observe CEνNS at the Spallation Neutron Source

Section 2.1: Challenges associated with a CEνNS observation

As discussed in Section 1.1, Freedman made some initial observations about the notable challenges that would be associated with any attempt to measure the CEνNS process, including: conspicuous observable in the form of a low-energy nuclear recoil; low interaction rates; and neutron backgrounds [111].

Challenging observable

The only signature of the CEνNS process is a low-energy nuclear recoil. Detection of nuclear recoils is made more challenging by the *quenching* of signal from energy depositions by heavier particles: the signal that is observed from a traveling proton, or alpha particle, or xenon nucleus, losing energy in the detector medium, is lower than the signal that would be observed if an electron were to deposit the same amount of energy. With quenching, the observed signal $I_{\rm obs}$ for a recoiling nucleus of energy $E_{\rm nr}$ could be expressed as

$$I_{\rm obs} = {\rm QF} \times Y \times E_{\rm nr},$$

where Y is the signal yield per unit deposited *electron equivalent* energy and QF is the quenching factor. Typically, the quenching factor is energy dependent and can be thought of as the fraction of deposited energy "available" for detection in a specified channel but the consequences of quenching are more subtle than simply making the signals lower amplitude. If, for instance, the quenched signal is near the threshold for detection, then the number of expected or observed events in an experiment can be especially sensitive to any uncertainty in the quenching factor (§3.2). Section 3.1.2 discusses some of the mechanisms by which quenching takes place.

Low interaction rates

As a process of the weak nuclear force, CEνNS cross sections are small compared to many other physical processes. Bethe and Peierls, having performed one of the early calculations of the inverse beta decay cross section and, finding it to be $< 10^{-44}$ cm^2, point out that this "[corresponds] to a penetrating power of 10^{16} km in solid matter" [48]. They go on to say that it is "therefore absolutely impossible to observe processes of

this kind" and close their letter with: "one can conclude that there is no practically possible way of observing the neutrino" [48].

Early cross-section calculations, such as those of Ref. [48], were generally low by a factor of two [26], but the commentary offered by Bethe and Peierls underscores just how fleetingly neutrinos interact with matter. Though a hallmark feature of CEνNS is a coherency-enhanced cross section, the increase is "only" an order of magnitude or two, leaving count rates low for even idealized experiments.

Backgrounds

Any detection system which is sufficiently sensitive to observe the low-energy nuclear recoils of the CEνNS process is necessarily quite sensitive to background radiation. Neutron backgrounds are particularly threatening: elastic neutron scattering results in the low-energy recoil of a nucleus, exactly as in the case of a CEνNS interaction. Consequently, even detection technologies laudably able to distinguish between electronic and nuclear recoils are still unable to reject neutron backgrounds without also disposing of CEνNS events. The significance of the challenge posed by backgrounds to the low-threshold astroparticle physics community is apparent in the effort invested to mitigate backgrounds [110].

2.1.1: Requirements for a successful CEνNS experiment

Generally and simplistically, a successful measurement of the CEνNS process must be built upon three pillars[1]:

Low-threshold detectors Even with a very-high total rate of CEνNS events in a detector, the rate of *detectable* events may remain meager and depends on the realizable threshold of the detector. Note that an additional "knob to turn" in detector selection is the target nucleus/nuclei: higher-mass targets benefit from increased event rates due to the N^2-like scaling of the cross section; however, these same heavier nuclei yield smaller signals due to decreased recoil energy for a fixed-energy neutrino source.

Source of neutrinos An abundant source of appropriate-energy neutrinos is necessary. Nuclear reactors present a prolific source of (anti)neutrinos, but the energies are generally $\mathcal{O}(1\text{ MeV})$, increasing the importance of a low threshold. Stopped-pion sources, such as the Spallation Neutron Source of Oak Ridge National Lab, the Los Alamos Meson Physics Facility (LAMPF[2]) at Los Alamos National Lab,

[1]These "pillars" are meant to serve as conceptual guidelines and do not represent orthogonal bases of the experimental-requirement space, which is to say that poor performance in one space could be compensated for by exceptional performance in another; for example: high background rates could be tolerable if the detector threshold and properties of the neutrino source combine to yield a manageable signal-to-background ratio.

[2]Now the Los Alamos Neutron Science Center (LANSCE).

or the ISIS facility at the Rutherford Appleton Lab present more energetic neutrinos that are generally suited for CEνNS searches but do not enjoy the same flux as reactors.

Low-background environment As an interaction moderated by the weak nuclear force, CEνNS event rates will be low, necessitating a low rate of background events in the signal region of interest (ROI). Since the CEνNS signal is expected to be contained at low energies, the overall background budget is strict. Properties of the neutrino source can also influence the background environment: for example, pulsed operation, possible with beam-driven sources, reduces susceptibility to steady-state backgrounds.

2.1.2: Sources of neutrinos

There are numerous sources of neutrinos that could be considered for a CEνNS search, but discussion here is restricted to the two most readily applicable to first generation CEνNS detectors: nuclear reactors and stopped-pion facilities.

Nuclear reactors are environments where an arrangement of actinide-based "fuel" and neutron-moderating control elements are configured in such a way as to maintain criticality in a chain reaction of nuclear fission. Fission yields daughter nuclei which are typically unstable with the net result being a very high rate of β-decays taking place in any given reactor core. Reactor neutrinos ushered into existence the field of experimental neutrino physics in 1956 with the confirmation of the very existence of the neutrino by Cowan *et al.* [78] and have continued to be a fount of insight.

As the ultimate source of reactor neutrinos is the β-decay of heavy radioactive nuclei, the yield is actually in the form of electron antineutrinos in the few MeV range. While this is relatively near energies that might be expected in supernovae, and thus might be particularly useful in attempts to measure astrophysically relevant neutrino interaction cross sections, this will restrict the nuclear recoil energies produced in CEνNS detectors and demand low-threshold capabilities.

Stopped-pion sources offer higher energy neutrinos but will generally require sacrifice in total flux. Also commonly referred to as decay-at-rest (DAR) sources, these provide a notable flux of ν_μ, ν_e, $\overline{\nu}_\mu$ neutrinos with kinematically well-defined energy distributions (§2.5.1). As an advantage, the energies of the neutrinos lend themselves well to CEνNS experiments: the 52.8 MeV maximum energy still maintains coherence and much of the energy distribution lies at energies more easily detectable than the population from nuclear reactors. Since the neutrinos at these facilities are ultimately produced by a particle beam, pulsed operation may be possible. One particular facility which can claim these two capabilities is the Spallation Neutron Source (SNS) of Oak Ridge National Lab (ORNL). The features of the SNS, and stopped-pion sources in

general, are discussed in greater detail in Secs. 2.4 and 2.5.

Section 2.2: Past proposals to measure CEνNS and dark-matter lemonade

The initial description of the CEνNS process included the first proposals for its measurement, invoking the possibility of "deuterium and helium bubble chambers, mineral oil or liquid helium scintillator tanks, and helium and neon streamer chambers" [111]. As the process would offer another check of the Standard Model, in addition to a wealth of other physics opportunities (see Secs. 1.3 – 1.6), there have been diverse, boundary-pushing proposals for its measurement from the experimental community in the decades since Freedman's first suggestion.

One of the early suggestions for detection of CEνNS signals came from Drukier and Stodolsky [94] who proposed using carefully configured superconducting grains to offer detectable signals from the low-energy nuclear recoils induced by CEνNS. In their proposed experiment, small $\mathcal{O}(1\ \mu\text{m})$ grains of a superconductor would be held in a magnetic field and at a temperature such that small changes in the temperature of the grain would cause it to "flip" from a superconducting state to a normal state; this flip results in a change of the nearby magnetic field which could be detected [94]. The detection scheme of Ref. [94] can provide fast timing information, but the binary nature of detection in each grain brings with it a disadvantage: any single grain either flips from the metastable, superconducting state, or it doesn't; the transition and detectable signal do not readily communicate any energy information beyond the exceeding of a threshold.

Drukier and Stodolsky's proposal resulted in recognition of a connection between CEνNS detection and another field: Goodman and Witten quickly identified that the superconducting grains of Ref. [94], capable of sensing low-energy nuclear recoils from CEνNS, would be equally capable of sensing low-energy nuclear recoils that resulted from interactions between dark matter candidates and nuclei [124]. The common modality of energy deposition between CEνNS and WIMP-like dark matter would prove mutually beneficial for experimentalists on both sides as any technological developments, or accrued experience and expertise, could be readily transferred between the fields.

Cabrera *et al.* [66] proposed bolometric detection of CEνNS soon after the appearance of the ideas of Ref. [94] and representing a somewhat more general approach. The connection proposed by Ref. [124] between CEνNS and WIMP searches would, when combined with the detection methodology of Cabrera *et al.* [66], started a recurring motif of CEνNS detectors becoming, instead, dark matter experiments. Bolometric detection of WIMPs, following the spirit of Ref. [66], is at the core of both the CDMS [73] and CRESST [74] dark matter detection efforts. Based on a different mechanism of detection, looking for ionization resulting from low-energy nuclear recoils rather than temperature changes, the development of p-type point contact

(PPC) germanium detectors was motivated largely by the possibility for CEνNS detection [34], quickly leading to the development of the Coherent Germanium Neutrino Technology (CoGeNT) experiment, which would become, again, a dark matter experiment [1].

Section 2.3: COHERENT at the SNS

In 2013, researchers from the fields of neutrino, dark-matter, and nuclear physics united with significant resolve to form the COHERENT Collaboration and make an unambiguous observation of the CEνNS process at the Spallation Neutron Source, ultimately using numerous detector technologies and nuclear targets. The connection between CEνNS and dark-matter detection technologies, underscored by the fact that interest in CEνNS detection led to the conception of numerous dark-matter searches (§2.2), would prove a boon for CEνNS efforts within COHERENT: numerous members had gained considerable experience with extremely sensitive experiments while searching for WIMPs. Other COHERENT collaborators brought extensive experience in neutron detection or sophisticated simulation and data analysis. Collectively, members of COHERENT represented all of the developments in nuclear, high-energy, and particle physics, and many members of the community that had long been proposing CEνNS measurements participated.

The SNS itself has been the subject of much interest within the neutrino community and those seeking CEνNS detection. Proposed experiments such as CLEAR [222] did considerable due-dilligence on the neutrino source and environment of the SNS. Despite years of sustained interested, neutrinos from the SNS were wasted until the COHERENT Collaboration was formed in 2013 and successfully established a foothold at the facility.

Section 2.4: The Spallation Neutron Source of Oak Ridge National Laboratory

The Spallation Neutron Source (SNS) is a billion-dollar user facility located at Oak Ridge National Laboratory (ORNL) designed to yield the most intense source of pulsed neutrons in the world, enabling neutron-scattering based research of materials, touching the fields of physics, chemistry, and biology [91, 179]. At the SNS, a liquid mercury target is struck by $\sim$1-GeV protons and the neutrons yielded by the resulting spallation reaction are utilized by experiments located at specialized target stations (24 possible target stations exist). The proton beam is delivered to the mercury target after acceleration in an initial linear accelerator and subsequent collection into physically-condensed charge "micro-bunches" which travel around an "accumulator" ring 248 meters in circumference [54]. As the beam is extracted from the ring and proceeds towards the target, the micro-bunches form a "mini" bunch which is distributed in time such that most of the charge is incident on the target within a 1-μs window.

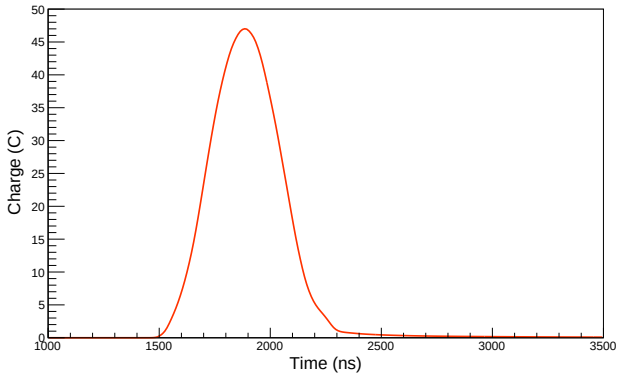

Figure 2.1: Average timing distribution of Spallation Neutron Source protons on target (POT). This represents POT timing profiles averaged over the period during which the data analyzed in Chapter 6 was collected, but represents well the shape expected for any individual pulse. Individual beam traces are measured by an inductive pickup coil after the proton bunch has been kicked from the accumulator ring and is approaching the liquid mercury target [151]. Averaging of beam pulses carried out by S.C. Hedges.

2.4.1: Beam timing characteristics

At various points in the acceleration system, the SNS proton beam passes through SNS beam current monitors (BCM), each composed of an inductive sensing loop and specially designed electronics [151]; signals from the BCM system are digitized at 100 MHz [68]. An average beam trace, recorded by the BCM closest to the spallation target and representing the time period over which data was collected for the analysis of Chap. 6, can be seen in Fig. 2.1.

Section 2.5: The SNS as the Spallation *Neutrino* Source

2.5.1: Neutrino production by stopped pions

When the high-energy proton beam strikes the mercury target, a large number of charged pions are produced. A considerable fraction of the π^+ population is moderated by the mercury target and subsequently decays at rest [197]. Decay of the stopped π^+, which have a lifetime of 26 ns, yield μ^+, lifetime 2.2 μs, which similarly thermalize and decay at rest. In addition to the μ^+, the positive pion decay produces a muon neutrino ν_μ, while the μ^+ decay produces populations of $\overline{\nu}_\mu$ and ν_e. The neutrinos produced in these processes have well-understood energy and time distributions which are discussed in subsequent sections.

The LSND [8] and KARMEN [23] experiments relied on similarly produced neutrino beams and address the issue of possible $\overline{\nu}_e$ contamination. Any π^- that are produced in the target moderate quickly and are captured by nuclei with only $\sim 1\%$ decaying prior to capture, producing μ^-. Of the μ^-, a high percentage will again be captured and fail to produce $\overline{\nu}_e$. The combination of these effects results in a negligible expected contamination by $\overline{\nu}_e$ well below the one-percent level [8, 23].

2.5.2: Anticipated neutrino spectra, timing, and flux from the SNS

The muon neutrinos produced in the decay of the moderated pions are monoenergetic with an energy given by [186]

$$f_{\nu_\mu}\left(E_{\nu_\mu}\right) = \delta\left(E_{\nu_\mu} - E_\pi\right), \qquad \text{with } E_\pi = \frac{m_\pi^2 - m_\mu^2}{2m_\pi},$$
$$\rightarrow E_{\nu_\mu} \approx 29.8 \text{ MeV}. \tag{2.1}$$

The $\overline{\nu}_\mu$ and ν_e produced from the decay of muons have well-defined energy distributions with maximum energies of $\sim$52.8 MeV; the energies of these neutrinos follow the Michel spectrum, with

$$f_{\nu_e}(E_{\nu_e}) \equiv \frac{dn_{\nu_e}}{dE_{\nu_e}} = \frac{96}{m_\mu^4}\left(m_\mu E_{\nu_e}^2 - 2E_{\nu_e}^3\right) dE_{\nu_e}, \tag{2.2a}$$

describing the distribution of ν_e energies and

$$f_{\overline{\nu}_\mu}\left(E_{\overline{\nu}_\mu}\right) \equiv \frac{dn_{\overline{\nu}_\mu}}{dE_{\overline{\nu}_\mu}} = \frac{16}{m_\mu^4}\left(3m_\mu E_{\overline{\nu}_\mu}^2 - 4E_{\overline{\nu}_\mu}^3\right) dE_{\overline{\nu}_\mu}, \tag{2.2b}$$

describing those of the $\overline{\nu}_\mu$ [19, 186]. In reality, there is a component of the SNS neutrino flux which is produced by leptons which decay prior to thermalization, yielding "decay-in-flight" neutrinos. A high-fidelity GEANT4 simulation of the neutrino production and transport at the SNS, including passage of the proton beam through the mercury target and the spallation process itself, yields the neutrino energy spectra shown in Fig. 2.2b [213].

The pulsed nature of the SNS presents considerable background-reduction opportunities, but realization depends on an understanding of the distribution in time of the neutrino pulse. Neutrino production times can be seen in Fig. 2.3.

Expected neutrino flux

The neutrino flux from the SNS can be expressed in terms of the number of decay-at-rest neutrinos produced per proton delivered to the spallation target. Though units of neutrinos per square centimeter per second may be a more natural way of expressing the flux, the realities of source operation are better reflected in terms of unit integrated beam power, leading to an expression for total produced neutrinos

$$n_\nu = \frac{P}{E_p} Y_\nu, \tag{2.3}$$

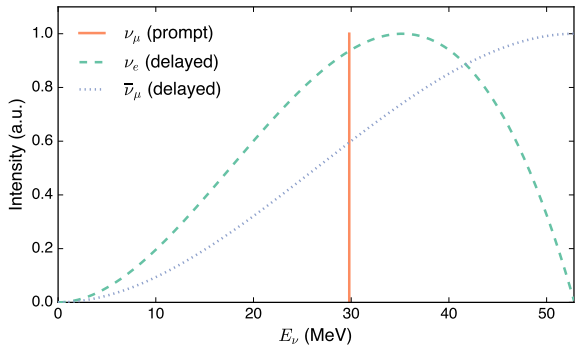

(a) Idealized, analytical energy distributions for the neutrinos produced at the SNS in the spallation process. The solid, coral line shows the monoenergetic ν_μ energy of 29.8 MeV; the dashed, aquamarine line represents the ν_e energy distribution; and the dotted, steel-blue line depicts the distribution of $\overline{\nu}_\mu$. The maximum neutrino energy is 52.8 MeV. Production of these neutrinos occurs with different timing characteristics (see text and Figure 2.3); the prompt neutrino population is composed entirely of ν_μ while ν_e and $\overline{\nu}_\mu$ constitute the delayed population.

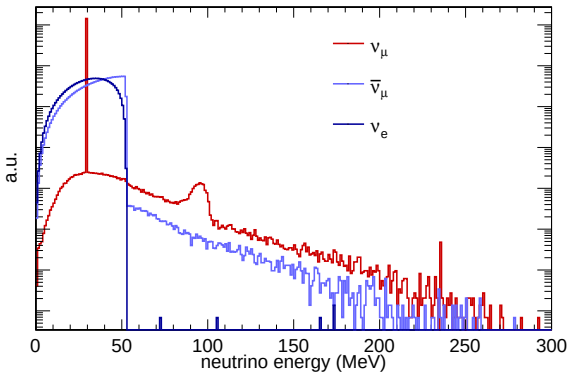

(b) Neutrino energy spectra expected from the SNS. These are the result of a GEANT4 simulation which includes the geometry of the SNS target and target building. The simulation includes transport of the SNS proton beam through the liquid mercury target, models the spallation process, and propagates the neutrinos through the building and to an area in the basement housing neutrino experiments. From The COHERENT Collaboration [17].

Figure 2.2: Energy distributions for the dominant neutrino species produced in the spallation process at the SNS. The left panel shows the idealized distribution, with monoenergetic ν_μ and the Michel spectrum, Equations (2.2), governing ν_e and $\overline{\nu}_\mu$. The right panel shows the neutrino energy distribution from simulation: decay of π^+ and μ^+ in flight yields much of the continuum; the feature at $\sim$100 MeV is associated with capture of μ^- [213].

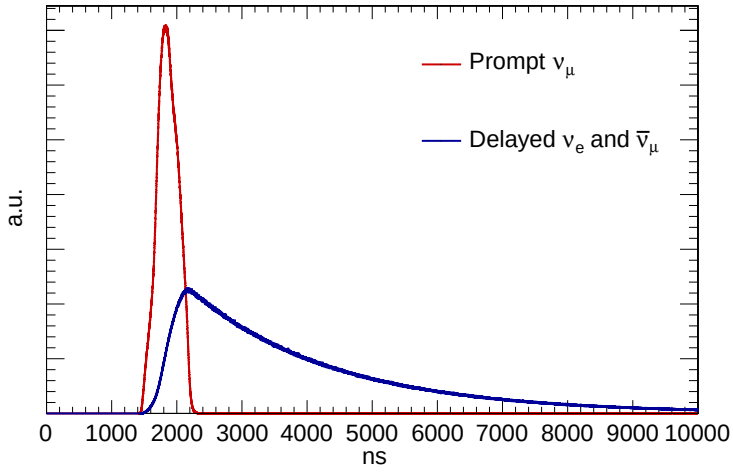

Figure 2.3: Anticipated distribution of production times for different neutrino species at the SNS. The prompt neutrino population, consisting entirely of ν_μ from pion decay, closely follows the timing of the protons on target (Figure 2.1). The delayed component consisting of equal parts ν_e and $\overline{\nu}_\mu$ results from decay of the muons which are, themselves, produced by pion decay process associated with the prompt neutrino population. These distributions are the result of: a convolution of the POT distribution like that shown in Figure 2.1 with an exponential of $\tau = 26$ ns (prompt component); and a convolution of the prompt component distribution with an exponential of $\tau = 2.2$ μs. Unlike the neutrino energy distributions seen in Figure 2.2, the timing distributions do not include decay-in-flight components. From The COHERENT Collaboration [17].

where P is the integrated beam power in units of GW $\cdot$ hr, E_p is the proton beam energy, and Y_ν is the number of ν produced per proton. Generally, Eq. (2.3) is flavor dependent. Simulation results from members of the COHERENT Collaboration found that, for typical SNS operating parameters as of 2015 – 2017 [213],

$$Y_\nu = \begin{cases} 0.08389 & \nu = \nu_\mu, \\ 0.08389 & \nu = \overline{\nu}_\mu, \\ 0.08368 & \nu = \nu_e. \end{cases}$$

Ultimately, a conservative 10% overall uncertainty on neutrino flux would be adopted by the COHERENT Collaboration, and a flavor-independent value of $Y_\nu \sim 0.08$ ν/p during typical operation could be used [17].

Section 2.6: COHERENT pathway towards unambiguous observation of CEνNS

Though the existence of the CEνNS process is non-controversial, the difficulty of producing an observation demands that considerable evidence be established that any observed signal is in fact associated with CEνNS. As an experiment with highly sensitive detectors, general background-reduction techniques will be applied and expertise from the rare-event search community will be employed [132]; additionally, certain location-

specific backgrounds must also be accounted for, such as those associated with neutrons from the SNS beam.

Beyond background considerations, there are a few key features of the CEνNS measurements of the COHERENT Collaboration which will reduce any ambiguities that could be associated with CEνNS observation. These features are:

Beam related The pulsed nature of the SNS presents the ability to establish with great confidence that a signal is associated with the SNS beam. The beam pulses are sufficiently well-contained in time with respect to the pulse period (see Section 2.4.1) that "coincident" and "anti-coincident" regions can easily be delineated.

Neutrino related Within a "coincident" region, where signals can already be associated with the SNS beam, it is possible to further determine (on a statistical basis) if a signal is likely to be associated with the neutrino component of the beam. The delayed neutrino population, produced from decay of muons, will follow a timing distribution unique from other beam features; observation of a muon-decay-like timing feature will thus provide confidence that features are not only related to the SNS beam but are associated with the neutrinos.

N^2 cross-section scaling Through the use of detector systems which are based on different target nuclei, the neutrino-related events can be shown to follow an N^2 shape. This dependence is unique from any expected backgrounds and will provide another layer of confidence that any signal excess is consistent with the CEνNS process.

The final component, observation of the $\sigma \propto N^2$ behavior, could be viewed as an especially rigorous requirement but its origin is in the quantum mechanics of the coherency of the CEνNS process and is thus a highly specific test. Though it is extremely unlikely that a non-CEνNS, neutrino-induced event would demonstrate an interaction probability similar to that expected for CEνNS, the N^2 behavior would retire any lingering ambiguity.

Section 2.7: Pioneering CEνNS detection effort with CsI[Na] and experiment siting

Initial detector deployments within the COHERENT Collaboration were focused on establishing an understanding of the backgrounds present in experimental area; some of these efforts are discussed in Sec. 5.5. The first CEνNS-detector selected for deployment was a sodium-doped CsI scintillator detector. Cesium iodide with a sodium dopant offers numerous attractive features as a CEνNS detector, as laid out by Collar *et al.* [77] and summarized below.

- Both Cs and I are both well-approximated as monoisotopic: ^{133}Cs and ^{127}I are the only stable isotopes for the two elements. This allows some simplification of the cross-section calculation as only a single configuration must be considered for each nuclear species in the CsI target. Multiple isotopic constituents would require independent calculation of the cross section for each isotope, as the overall vector and axial coupling constants, G_V and G_A given by Eqs. (1.2a) and (1.2b), respectively, are dependent upon nuclear configuration.

- Cesium and iodine are both relatively heavy nuclei with correspondingly high numbers of neutrons; for Cs, $N = 78$, and for I, $N = 74$. This is an advantage for CEνNS detection as the N^2 cross-section scaling (§1.2) thus affords an enhancement to the number of CEνNS events in a CsI target over lighter-mass (lower-N) targets.

- The masses of Cs and I are sufficiently similar as to be approximated by the same effective mass. This allows the same quenching factor to be used for both recoiling species.

- The sodium-doped CsI scintillator does not demonstrate the same long-timescale afterglow seen in the thallium-doped variant, mitigating the impact of background radiation sources which might otherwise yield scintillation photons long after interaction with the crystal, contaminating potential CEνNS events [154].

- Low-background CsI crystals can be commercially produced [81] and are relatively inexpensive, $\sim$ \$1/g [77].

An initial plan for the deployment was described by Collar *et al.* and involved potential installation of the detector system underground [77]. Underground installation is *de rigueur* for direct dark-matter detection efforts, as the overburden attenuates the exposure to cosmic-ray backgrounds, which can be crucial to obtain the best performance from highly sensitive detectors [110, 116, 132]. Exploration of the SNS facility and consultation with building supervisors presented an attractive alternative: a basement hallway.

Though seemingly inglorious, "neutrino alley", as the hallway came to be known, possessed several virtues which could help bring an observation of CEνNS to fruition. A modest $\sim$8 m.w.e. overburden[3] allows for some reduction of the cosmic background, which is already substantially reduced by the pulsed operation of the SNS. If hardware problems arise, or if calibrations need to carried out, experiments deployed in neutrino alley remain accessible to researchers. Finally, the hallway allows placement of experiments at relatively short distances from the spallation target itself, with the distance of closest approach $\sim$20 m.

[3]Meters water equivalent (m.w.e.) describes the depth of water that would be necessary to achieve the same reduction in background as realized by different materials. A detailed discussion can be found in Refs. [110, 132].

As the hallway is part of the existing SNS infrastructure, locations within the hallway can be determined to an accuracy much smaller than the distance to target, allowing the standoff uncertainty to be neglected. The SNS Survey and Alignment team has a high-fidelity LIDAR survey system which produces location "point clouds" that are fiducialized within the SNS facility based on previous scans. With this capability, a dedicated LIDAR scan of neutrino alley can be placed accurately with respect to other SNS features. Based on these scans, the CsI[Na] detector would ultimately be placed at a distance of 19.3 m from the neutrino source. The uncertainty on this standoff distance is on the order of centimeters, per the survey team, and is negligible compared to other sources of uncertainties in neutrino flux estimates (§2.5.2).

Information from these scans was also used to inform a design plan for neutrino alley, shown in Fig. 2.4. This drawing provides some context for the CsI[Na] placement, and neutrino alley, with respect to the SNS facility. Additional aspects of the COHERENT suite of experiments are shown in the drawing, as well: the Sandia camera and SciBATH constituted important components of the neutron background measurements carried out; the NIN cubes are intended to provide measurements of the neutrino-induced neutron process in various materials and are discussed in Sec. 7.8; and the NaI[Tl] and CENNS-10 detectors are intended to provide CEνNS observations on argon and sodium. A photograph taken during the installation of one of the COHERENT experiments can be seen in Fig. 2.5, looking down the hallway in a direction antiparallel to the proton beam in Fig. 2.4.

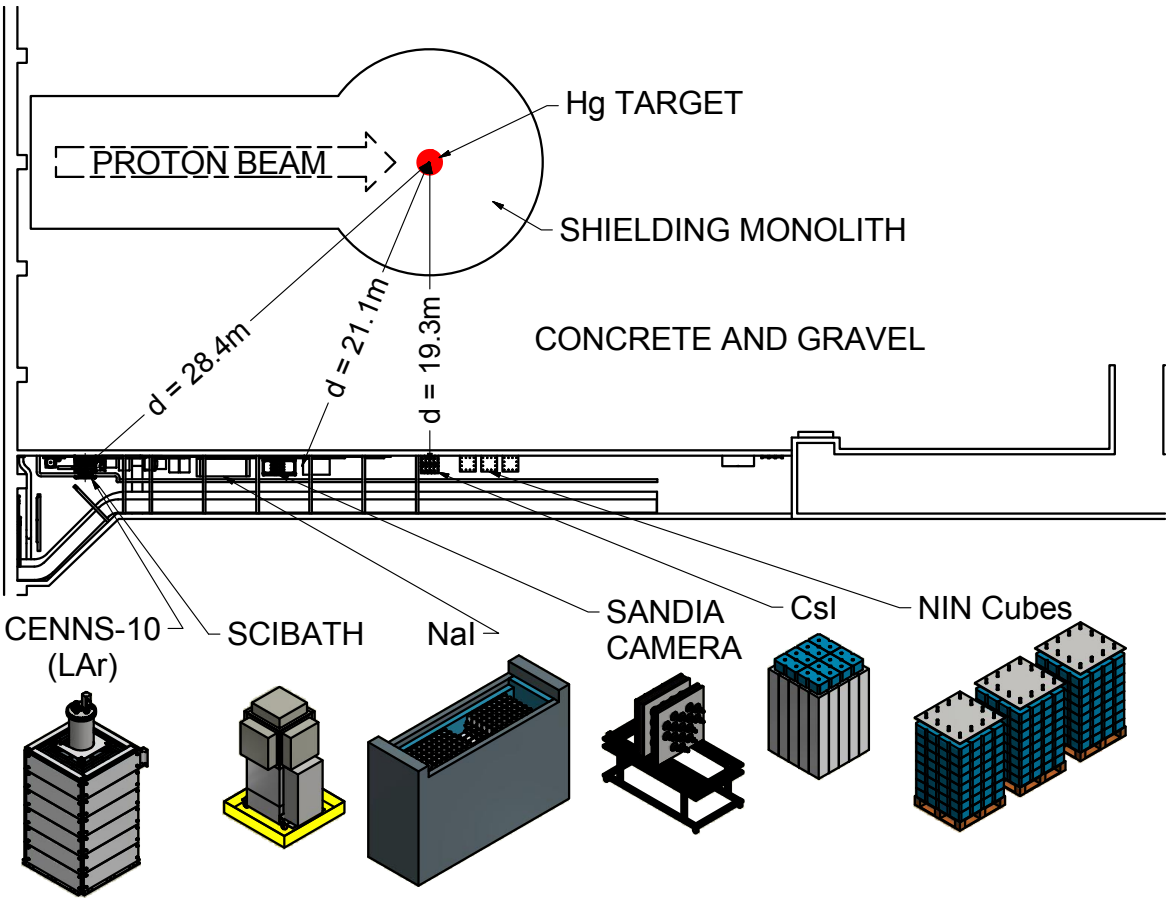

Figure 2.4: Siting of the detectors within the COHERENT effort along "neutrino alley" at the SNS. Representative distances from the center of the target are shown; these are based on a precise scan of the corridor carried out by the survey and alignment team of the SNS. From the COHERENT Collaboration [17].

Figure 2.5: An image of "neutrino alley", the location at the SNS where the COHERENT suite of experiments is located. The right side of the hallway is located closest to the SNS target. On the left side of the hallway, above the scaffolding, is the hot off-gas pipe referred to in Section 5.2.3. This photograph was taken during the installation of electronics for the measurement of the neutrino-induced neutron production process on lead, discussed in Section 7.8. The configuration at this time was somewhat different than that shown in Figure 2.4, but the CsI[Na] experiment and the lead neutrino cube are visible in their final locations in the distance on the right side of the hallway.

CHAPTER 3: Quenching factors

Section 3.1: Generation of detectable signal from CEνNS in CsI[Na]

3.1.1: Energy deposition by low-energy nuclear recoils

As has been discussed, a CEνNS interaction results in the recoil of a target nucleus with only modest kinetic energy. Prior to "recording" of this event, the energy of this recoiling nucleus must be transferred to a scintillation center; in CsI[Na], this transfer is generally facilitated by first transferring energy to the CsI crystal lattice [53]. The actual transfer of energy from a recoiling nucleus to the surrounding crystal lattice likely [17] takes place through mechanisms closely related to track formation in materials which is a combination [59] of Coulomb explosions [105] and the mechanism underlying the functioning of bubble chambers, thermal spikes [230].

In a Coulomb explosion, a region of material quickly becomes deficient in electron density due to the traveling ion; the newly positive ions in the region then mutually and strongly repel each other [105]. Thermal spikes pertain more directly to the motion of the atoms in a material and represent a scenario when local kinetic energies correspond to very high temperature values [215, 230], as might exist very briefly in a localized area around a (highly ionizing) nuclear recoil prior to dissipation of the kinetic energy through the material. Both effects are many-body problems, and standard stopping-power calculations in nuclear/particle physics which rely on the binary collision approximation (i.e., SRIM [258]) are inadequate to capture the full details, requiring molecular dynamics (MD) simulations [215]. Long thought to be separate, "competing" explanations for track formation [59, 105], molecular dynamics simulations have suggested that Coulomb explosions and thermal spikes both participate in track formation. The simulations of Ref. [59] point to track formation being a multi-stage process, with Coulomb explosions and thermal spikes representing the early- and late-time stages, respectively.

Irrespective of the precise process by which the energy is transferred to the lattice, the end result is the production of electron-hole pairs which can travel through the lattice. These pairs can then recombine to form an exciton, which continues to diffuse through the lattice; alternatively, they can be captured by traps or continue to diffuse separately until they recombine or trap on a luminescence center (Na) [53]. If excitons are produced, they diffuse until either trapped in a lattice vacancy or a quenching center or until they are

captured by a luminescence center [53]. For a full discussion of scintillation in inorganic crystals, see Ref. [53].

3.1.2: Generation of scintillation light and quenching

Cesium iodide is part of a broader class of inorganic crystalline scintillators which share common light-production mechanisms [53]. Though undoped CsI does scintillate, the discussion here will focus on the scintillation produced by the Na dopant.

Dopant centers are excited by either excitons or the separate capture of an electron and a hole and the scintillation light is then a consequence of deexcitation of the luminescent center to its ground state [154]. Several variations on this theme can give rise to observable characteristics. Excitation of a luminescent center to a metastable level, whose direct decay to the ground state is forbidden by photon emission, can give rise to phosphorescence: light emission over the course of a relatively long timescale, as the metastable dopants must be further excited, via lattice phonons for instance, to a level from which it can decay to its ground state [154]. Additionally, nonradiative transitions from excited states of the luminescent centers to the ground state are one way by which quenching is realized: the energy is ultimately dissipated in a way that does not yield a detectable photon.

Quenching in inorganic scintillators is also suspected to take place via "ionization quenching" [53]. This process was originally proposed as a feature of organic scintillators, wherein localized and temporary molecular damage arising from high ionization density (such as that which is created by heavy ions subject to high dE/dx) is "probably" the cause of nonradiative energy dissipation (i.e., quenching) even in the case of inorganic scintillators, though the mechanism is unexplained [53].

There are numerous proposed models for the behavior of quenching at low energies throughout the literature. Tretyak [242] offered a semi-empirical model whose theoretical underpinnings are attractive: the behavior is based on the traditional Birks' law behavior combined with stopping-power calculations provided by the widely-popular SRIM [258] and ESTAR [43] utilities. Birks' law [52] describes the light per deposited energy dS/dr with the equation

$$\frac{dS}{dr} = \frac{A}{1 + kB\left(\frac{dE}{dr}\right)}\frac{dE}{dr}, \tag{3.1}$$

where A and kB are constants, the latter referred to as the Birks factor. Though originally developed to describe the scintillation yield in organic crystals, Birks' formula, as Eq. (3.1) came to be known, has shown success at describing the scintillation yields for many materials, extending beyond the original organic scintillators for which the theory was developed in Ref. [52]. The proposal by Tretyak to base inorganic scintillator response models on a model developed for organic scintillators finds support in the evidence for

ionization quenching in inorganics [53], as this process is considered in the development of Eq. (3.1). Despite their significance and widespread use, inorganic crystalline scintillators are without a model which fully considers all of the microphysics involved in the light production process, and experimentalists seeking to model the response of these materials are left with little recourse than to use the semi-empirically motivated method of Ref. [242].

Section 3.2: Effect on threshold and efficiency uncertainties in CEνNS and similar experiments

Uncertainty on the QF can have considerable impact in low-threshold astroparticle physical measurements. The focus here is directly on the consequences for CEνNS measurements, but similar issues exist with, for instance, dark matter searches.

If the CEνNS search is considered a counting experiment alone, then the rate of CEνNS recoils above detection threshold is the quantity of interest. To quantitatively consider the impact of QF uncertainty, the recoil distributions for CEνNS shown in Fig. 7.1 are explored. Specifically, we focus on the impact of QF on a heavy target like CsI and a light target, like Na recoils in NaI[Tl].

Figure 3.1 shows the number of integrated counts "over threshold" from the curves of Fig. 7.1 as a function of the threshold in nuclear recoil energy. To perform the comparison, we assume a nominal detection threshold of 5 photoelectrons: the more "natural" unit in which threshold would be established and expressed. We assume a QF of 8.78% for CsI[Na] and 15% for Na recoils in NaI[Tl], in rough accord with the state of literature for CEνNS-appropriate recoil energies (see, e.g., Refs. [75, 257]). Further, the photoelectron yield used for CsI[Na] is 13.35 PE/keVee in accord with the detector used in the CEνNS search discussed subsequently; a light yield of 38 photons/keVee is adopted for NaI[Tl] [154] and a scaling factor of 0.2 approximates the quantum efficiency of a standard bialkali phototube, giving a photoelectron yield of 7.6 PE/keVee.

For both CsI[Na] and NaI[Tl], a 30% relative uncertainty is assumed on the QF. The 5 PE threshold is drawn along the recoil energy axis of Fig. 3.1 as a vertical solid red line; dashed vertical red lines illustrate the $\pm 1\sigma$ uncertainty band. Horizontal black lines extending from the integrated-count distributions to the vertical axis assist the eye in identifying the impact of the QF uncertainty on expected count rates; a solid horizontal black line is the mean expected count rate, while dotted lines again show the $\pm 1\sigma$ uncertainty band. In the given scenario, the 30% QF uncertainty results in a count rate uncertainty of $(-5.5\%, 2.8\%)$ for NaI[Tl], corresponding to the lower and upper boundaries of the uncertainty interval; for CsI[Na], the QF leads to a rate uncertainty of $(-32.9\%, 23.0\%)$.

Clearly, the relative importance of having a precision QF measurement depends somewhat on the mass

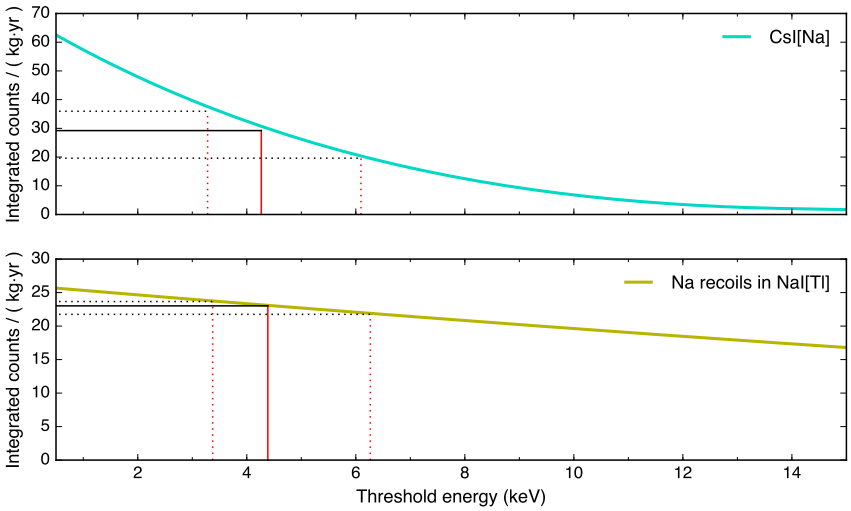

Figure 3.1: Demonstration of the effect of QF uncertainty on integrated count rate for CsI[Na] (top panel) and Na recoils in NaI[Tl] (bottom panel). The curves show the number of total integrated counts over threshold and the horizontal axis is the threshold. Recoil curves which generate the integrated-count distributions for both targets are generated by the CEνNS differential cross section, but the absolute scaling is arbitrary. QF values are material-appropriate but both are given a nominal relative uncertainty of 30%. The steepness of the recoil distribution for heavier nuclei (see Figure 7.1) results in an exaggerated impact of QF uncertainty on the number of counts above threshold, indicated by the horizontal black lines extending from the integrated count curves to the vertical axes; the solid black line indicates the mean integrated count rate expected, with $\pm 1\sigma$ indicated by the dotted black lines.

of the target nucleus. CsI[Na] happens to translate the QF relative uncertainty almost directly into the relative uncertainty on the expected count rate, but experiments with lighter nuclei may be more forgiving and reduce the impact of uncertainty on the detector response characteristics.

Section 3.3: Methods of measuring quenching factors

3.3.1: Direct ion irradiation

In the case of a measurement of the quenching associated with ^{4}He recoils in helium gas, the authors of Ref. [220] had a low-energy beam of α particles strike their sensitive volume. In this configuration, with common probe and target species, excitations of the medium due to the incident probe mimic the self-recoils generated by elastic scattering of particles in the detector. Even with similar probe and target nuclei, a possible issue with such an approach is differences associated with the ionization state of the recoiling particle. Alpha-particle beams have been used in instances to assess the quenching in other materials, such as liquid argon [165]. In the case of a CEνNS or WIMP search, the primary quantity of interest is self-recoils (i.e., argon nuclei recoiling in an argon medium). Combined especially with possible charge-state concerns, generalization from an incident beam of light ions to internally generated heavy-nucleus recoils is likely

nontrivial.

3.3.2: Nuclear resonance fluoresence

Joshi [147] recognized that the process of nuclear resonance fluorescence [153] could provide access to extremely low-energy nuclear recoils that result from momentum transfer by absorbed and reemitted photons. This technique is dependent on the identification of viable NRF states in the nuclear target, as the excited nucleus must reemit the NRF photon prior to loss of energy to the surrounding medium; this restricts useful states to those with very short lifetimes [147]. Additionally, the use of a photon beam significantly more broad than the target NRF resonance can result in a low signal-to-noise ratio, as non-NRF photon interactions contribute significantly to the total number of events in the detector. For relatively slow detector systems, insufficiently narrow bandwidth of the photon beam can result in a "blinded" detector, where a substantial number of the legitimate NRF events are contaminated by accidental coincidences with other beam-photon events.

The NRF-based QF measurement technique remains undemonstrated but could become more viable with the development of more brilliant, narrow-bandwidth γ-ray sources.

3.3.3: Neutron scattering

Quenching factors measurements can be complicated by the need to produce recoiling nuclei in a target medium while avoiding atomic or electronic excitations. The avoidance of ancillary excitations of the medium largely precludes the use of charged-particle probes, so the focus typically falls onto neutron-scattering experiments. Two complementary approaches for QF measurement using neutron scattering are described here. They are complementary in the sense that they have a degree of systematic independence: tagged scattering selects a narrow range of recoil energies for analysis, but makes the assumption that such a selection can cleanly and reliably be made; endpoint measurements, on the other hand, provide a measure of the QF at the maximum kinematically allowed recoil energy $E_{\rm nr}^{\rm max}$ for an incident neutron beam, but the data will contain events with a continuum of recoil energies extending from 0 to $E_{\rm nr}^{\rm max}$.

Nuclear recoils from neutron beams using energy endpoint

An incident neutron beam of a given energy, with endpoint $E_{\rm max}$, will produce nuclear recoils with energies ranging between 0 and [148]

$$E_{\rm nr}^{\rm max} = \frac{4E_{\rm max}m_{\rm nuc}m_n}{\left(m_{\rm nuc}+m_n\right)^2},$$

where m_{nuc} and m_n are the mass of the recoiling nucleus and the neutron, respectively. Any experiment which uses a continuous recoil distribution must be careful to account for this in the analysis, but this is something which can be addressed.

The neutron-recoil endpoint technique has been used, for instance, by Joshi *et al.* [148] to determine QFs in liquid argon, motivated by interest in using such a detector to measure CEνNS [219]. Recent work by Stiegler *et al.* [235] has demonstrated the use of this approach in inorganic crystal scintillators, obtaining a measurement of the QF for sodium recoils in NaI[Tl].

Production of nuclear recoils with tagged elastic neutron scattering

In the tagged-recoil approach, an elastically scattered neutron is detected and the angle of its scattering is determined based on parameters of its detection (i.e., it is found in a detector positioned at a well-known recoil angle). If the energy distribution of the incident neutron beam is known, then the distribution of nuclear recoil energies associated with elastic scattering into a known angle by a nucleus with a known mass is well defined. The energy transferred to the recoiling nucleus is given by [24]

$$\Delta E = \frac{2E_n M_n^2}{\left(M_n + M_T\right)^2}\left[\frac{M_T}{M_n} + \sin^2\theta - \cos\theta\sqrt{\left(\frac{M_T}{M_n}\right)^2 - \sin^2\theta}\right] \tag{3.2}$$

In cases where the recoiling target nucleus is much more massive than the incident neutron, $M_T \gg M_n$, this equation can be simplified to

$$\Delta E \approx \frac{2E_n M_n M_T}{\left(M_n + M_T\right)^2}\left(1 - \cos\theta\right). \tag{3.3}$$

Figure 3.2 shows a schematic of a quenching factor experiment based on elastic neutron scattering. Quasi-monoenergetic neutron beams and narrowly-defined recoil angles θ clearly improve precision that can be realized.

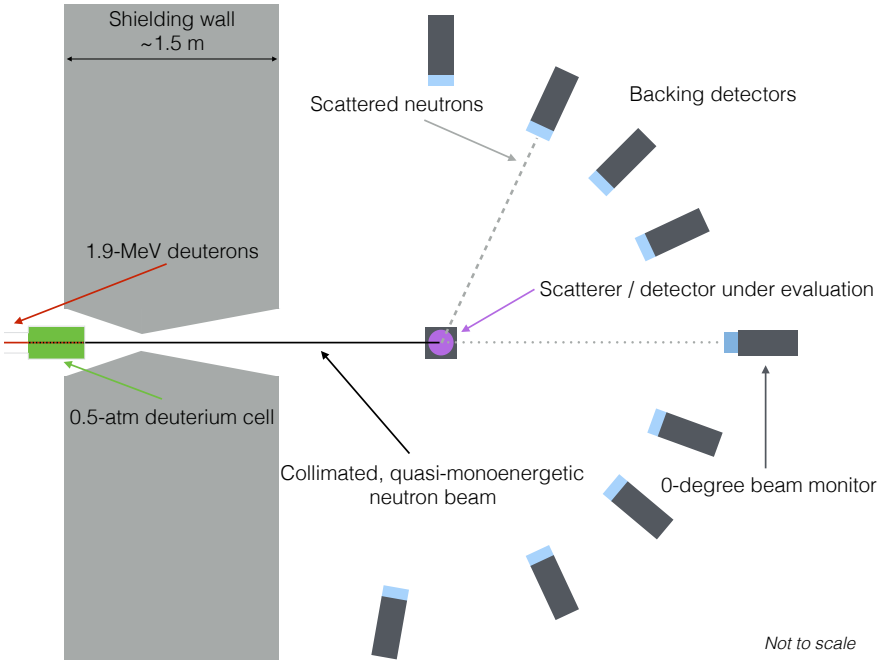

Figure 3.2: Schematic representation of a quenching factor measurement experiment utilizing neutron scattering. Neutrons are produced through the use of a charged-particle beam incident upon an appropriate nuclear target (at left). The neutron emission will generally be into a large solid angle, and depending upon the experimental parameters collimation of this source, as depicted, into a well-directed beam may be appropriate or necessary.

CHAPTER 4: Measurement of nuclear-recoil quenching factors in CsI[Na]

Section 4.1: Neutron beam production at TUNL

4.1.1: The TUNL tandem accelerator and the Shielded Source Area

Neutron beams for these measurements were produced at the tandem van de Graaff accelerator facility of TUNL, an overview of which is shown in Fig. 4.1. Here, we will briefly introduce aspects and capabilities of the lab relevant to the CsI[Na] QF measurement.

A beam of D^- ions was produced by a duoplasmatron source [166] biased relative to the subsequent beam-transport system so that the deuterons had an energy of 50 keV. The beam was acted on by a series of two electrostatic "choppers" followed by a "buncher", constituting a system which produces a pulsed ion beam. The pulsing system was driven by a single master oscillator which runs at 5 MHz but drives the system in such a way that the minimum pulsed-beam-delivery frequency is 2.5 MHz; by adjusting the oscillator output connected to one of the choppers, pulse arrival periods given by $t_{\text{pulse}} = 400 \times 2^n$, with n an integer ≥ 0, can be chosen, with transmitted beam current reduced by a factor of 2 for each successive n.

This pulsed, negatively-charged beam entered a tandem van de Graaff accelerator [246] from High Voltage Engineering Corporation, model FN, which has undergone conversion to a pelletron charging system [131] with a maximum terminal voltage of 10 MeV. Negative ions are accelerated towards the positively-charged central terminal of the accelerator where they pass through a thin carbon foil which strips the electrons [251], leaving a positive ion which undergoes a second stage of acceleration *away* from the terminal [246].

Following acceleration, the high-energy beam passes through an analyzing magnetic spectrometer. The field of this magnet, unlike other steering elements in the beam delivery system, was precisely monitored by an NMR probe whose measurements provide feedback to a control loop for the current delivered to magnet. Precise control of the field in this magnet allows for confident specification of the energy of the beam delivered downstream of the spectrometer; an extensively qualified calculator is available [252] to provide appropriate field settings for common beam species and experimental beam lines at TUNL.

Upon exit from the magnet, the beam passed through a series of "slits": a vertical aperture, centered in the beam line, defined by conducting fins from which the beam current can be measured. Since the magnetic field and the magnet geometry are well defined, an imbalance between the two "slit currents" indicates the

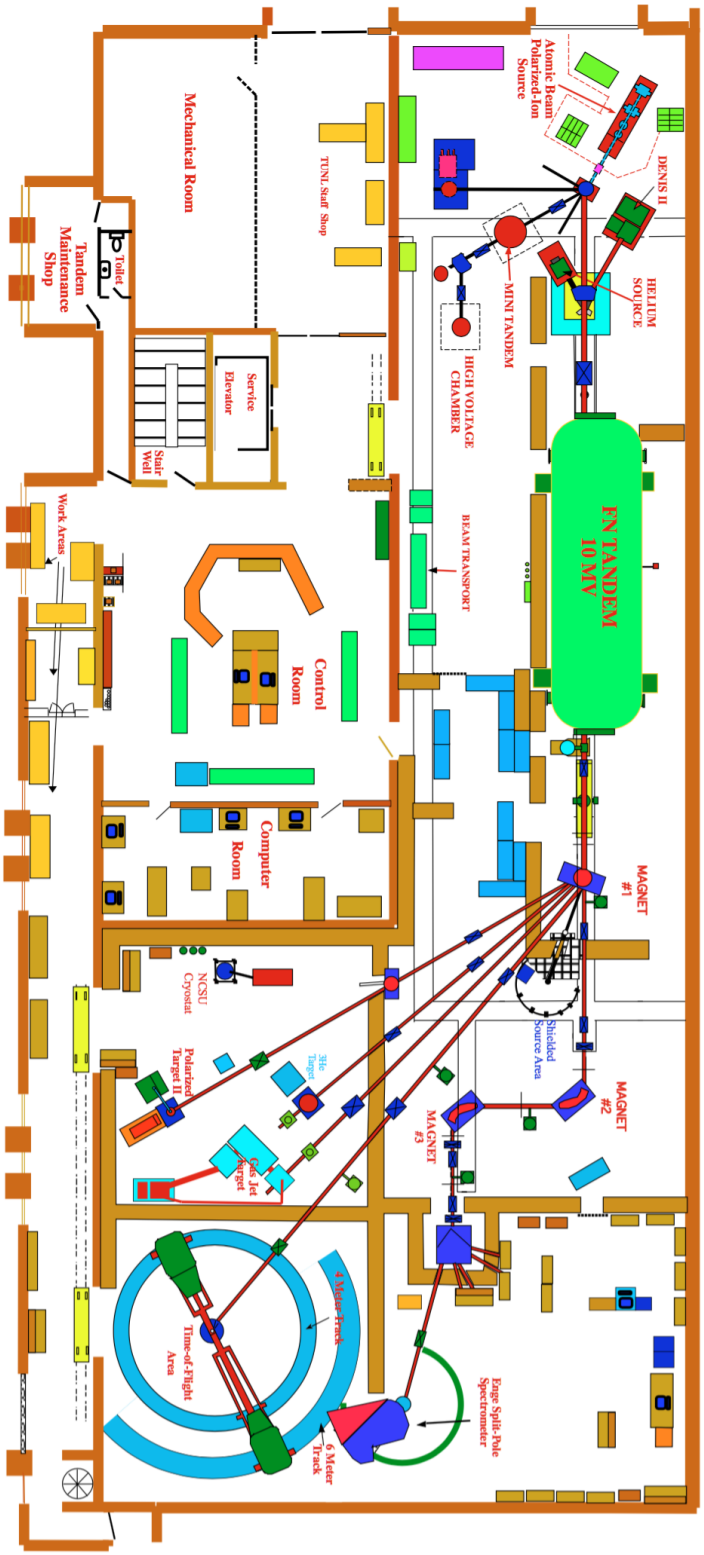

Figure 4.1: Overview of the TUNL tandem accelerator laboratory. The experimental area utilized for the measurements discussed in detail in Chapter 4, the Shielded Source Area (SSA), is visible in the lower-right quadrant of the diagram between "Magnet #1" and "Magnet #2". Image from A. Crowell.

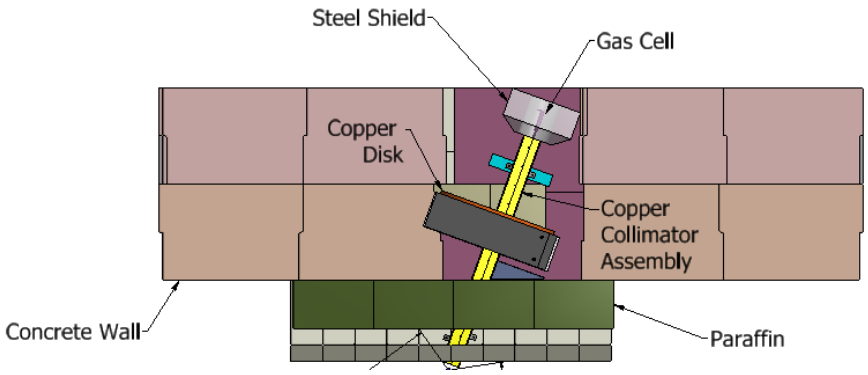

Figure 4.2: Top-down, cross-sectional view of the SSA shielding structure, designed to highly attenuate the transmission of off-axis beam neutrons. Neutrons are produced in a deuterium gas cell, positioned at the end of a charged-particle beam line. A copper collimator transmits neutrons emitted along the direction of the charged-particle beam to the experimental area of the SSA; the distance from the gas cell to the exit of the collimator is ∼1.5 m. Figure from M. Emamian.

ion beam is not of the desired energy defined by the magnet settings. Once sufficient current is detectable on the slits, a control loop is activated which adjusts the charge on the central terminal of the accelerator via regulation of the corona current drawn from the terminal in an effort to minimize the slit current. This feedback loop, once engaged, facilitates high-stability operation and very-precise beam-energy control.

For the purpose of this experiment, the deuteron beam was deflected 20° by the analyzing magnet along a beamline associated with the "Shielded Source Area" (SSA). The SSA was designed by TUNL researchers for measurements which rely on detection of scattered neutrons but are very sensitive to accidental backgrounds; to this end, the SSA has a sizable, purposefully designed shielding structure which admits only a tightly collimated neutron beam into the experimental area. Figure 4.2 shows a top-down cutaway of the SSA shielding structure while Fig. 4.3 shows the areal profile of a neutron beam in the SSA, measured for a separate experiment by Duke University graduate student Ron Malone. The qualitative features of Fig. 4.3 demonstrate the tight collimation of neutron beams which makes the SSA an especially attractive site for QF measurements like those described in Sec. 3.3.3 and undertaken here.

4.1.2: Neutron production using the $\mathrm{D(D},n)^3\mathrm{He}$ reaction

TUNL possesses the capability of producing neutron beams with a wide variety of energy-distribution characteristics, covering a wide range of energies, through bombardment of suitable targets with the accelerated charged-particle beams from the tandem accelerator. Many reactions can be used for neutron production, including, but certainly not limited to: $^7\mathrm{Li}(p,n)$, $\mathrm{D(D},n)^3\mathrm{He}$, $^3\mathrm{H}(p,n)^3\mathrm{He}$, and $^3\mathrm{H(D},n)^4\mathrm{He}$. The ranges over which any given neutron-production reaction is "useful" is actually somewhat specific to the experimental configuration: some reactions only yield monoenergetic neutron populations between specific

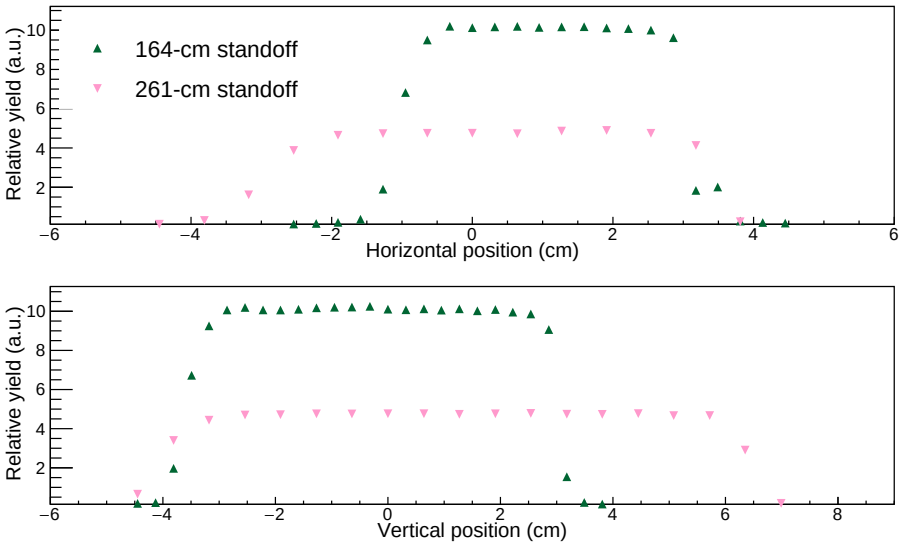

Figure 4.3: Example of SSA collimated neutron beam X-Y intensity measured at two distances from neutron-production cell. Even several meters into the experimental area of the SSA, the neutron beam remains well-contained within an area of several square inches. The data shown was collected by R. Malone of Duke University/TUNL for a separate measurement and reflects the qualitative properties of the neutron beam realized for the CsI[Na] QF measurement presented here.

energies. Experiments which can tolerate broad energy distributions or distributions with multiple energy populations may find different reactions more advantageous for use than an experiment which relies on a well-defined, narrow neutron energy distribution. For the experiment here, the $\mathrm{D}(\mathrm{D},n)^3\mathrm{He}$ reaction is selected for its high cross section and production of quasimonoenergetic neutrons within the desired energy range.

The $\mathrm{D}(\mathrm{D},n)^3\mathrm{He}$ reaction is exothermic with a positive Q value of 3268.904 keV [206, 248] and thus, aside from the Coulomb barrier arising from two like-charged nuclei repelling each other, has no threshold and the products are reasonably energetic irrespective of the bombarding energy; in this case, the lowest energy neutrons are produced with energies of ~ 2.5 MeV. Ample experimental data on this reaction exists and is well summarized by Liskien and Paulsen [173], who have specifically tabulated recommended values of the differential cross section at zero degrees, the total cross section, and Legendre Coefficients for a range of incident deuteron energies.

The energy of the produced neutrons is precisely determined by kinematics with negligible relativistic effects for the energies of interest to these investigations. Considering, in the lab reference frame, a deuteron with kinetic energy E_d incident on another deuteron at rest, we have from Ref. [141],

$$\sqrt{E_n} = r \pm \sqrt{r^2 + s}, \tag{4.1a}$$

with

$$r \equiv \frac{\sqrt{m_d m_n E_d}}{m_n m_{^3\mathrm{He}}} \cos\theta_n, \tag{4.1b}$$

and

$$s \equiv \frac{E_d\,(m_{^3\mathrm{He}} - m_d) + m_{^3\mathrm{He}} Q}{m_n + m_{^3\mathrm{He}}}, \tag{4.1c}$$

where θ_n is the lab angle of the outgoing neutron relative to the momentum vector of the incident deuteron and $m_{n,d,^3\mathrm{He}}$ refer to the masses of the neutron, deuteron, and helium-3 nucleus, respectively. Requiring real solutions for E_n in Eq. (4.1a), the solution is single valued, accepting only the '+' form. Neutrons produced along the incident deuteron-beam axis (i.e., at zero degrees; $\cos\theta_n = 1$) thus have an energy given by

$$E_n = \left[\frac{\sqrt{m_d m_n E_d}}{m_n + m_{^3\mathrm{He}}} + \sqrt{\frac{m_d m_n E_d}{(m_n + m_{^3\mathrm{He}})^2} + \frac{E_d\,(m_{^3\mathrm{He}} - m_d) + m_{^3\mathrm{He}} Q}{m_n + m_{^3\mathrm{He}}}}\right]^2. \tag{4.2}$$

For the measurements discussed in this work, a small deuterium gas cell was kept at a pressure of ∼0.5 atm; the below-atmosphere pressure was chosen to limit the amount of energy loss of the deuteron beam through the cell, correspondingly reducing the range of neutron energies produced. Figure 4.4 shows a schematic of the target cell configuration. The "target" region is the tip of this assembly, isolated from the charged-particle beam line, which is held at vacuum, by a thin (∼6.35 μm) havar window [79]. A thin tantalum sleeve and disk are inserted into the tip of the cell; these inserts are sufficiently thick to stop $\mathcal{O}$(1MeV) deuterons before they are incident upon the copper structural walls of the cell, mitigating the number of nuclear reactions likely to be induced by the beam. The gas-filled region of the cell has a diameter of ∼0.8 cm and a length of ∼2.8 cm.

4.1.3: Modeling of D(D,n)^{3}He neutron beams

A precise measurement of the quenching factor from neutron scattering depends on a precise knowledge of the incident-neutron energy distribution (see Eq. (3.2)). The use of a pulsed charged-particle beam affords the opportunity to carry out measurements of the neutron beam energy via time-of-flight techniques. A numerical model for the TOF spectra, accounting for finite geometries, energy loss, energy-dependent cross sections, and other issues is described in Sec. D.4. Extraction of energy information from the collected data is discussed more generally through the entirety of App. D.

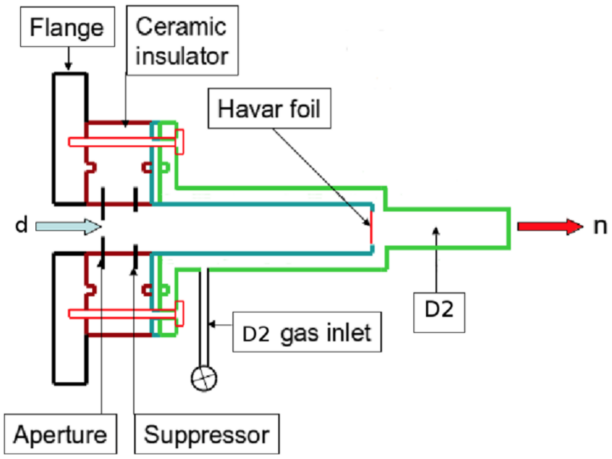

Figure 4.4: Schematic of the deuterium gas cell used to produce neutrons in the SSA. The outer, structural material of the gas cell is copper. A thin tantalum sleeve and disk are inserted into the tip of the cell, stopping beam deuterons and reducing the likelihood of beam-induced reactions on the copper structure. Figure from Reference [51].

Section 4.2: Experimental setup

4.2.1: Physical installation in the shielded-source area at TUNL

Over a period of 2 weeks during January and February of 2016, data was collected at the Shielded Source Area (SSA) of the TUNL tandem accelerator laboratory. The SSA is discussed in Sec. 4.1.1 and its location in the context of the TUNL tandem lab can be seen in Fig. 4.1.

Twelve 2"-diameter, 2"-long EJ-309 liquid scintillator detectors[1] were used as **backing detectors**: positioned around, and facing, the central scattering target being characterized but outside of direct illumination by the neutron beam, each of these would "tag" events where a beam neutron is scattered into the solid angle subtended by the detector (see Fig. 3.2). The neutron beam was directly monitored by a 2"-diameter, 1.5"-thick cell of BC-501A liquid scintillator [80].

The QF measurement was set up on top of an aluminum, semi-circular table and the 12 backing detectors were positioned at 11 unique angles. Table 4.1 shows the standoff distance from the scattering detector and the recoil angle tagged by each of the backing detectors. A photograph of the experimental setup can be seen in Fig. 4.5.

The CsI[Na] crystal assembly was acquired from Proteus [142] and generously loaned to TUNL by Prof. Juan Collar of the University of Chicago for the purposes of these measurements. The crystal itself, a right circular cylinder, measured 19 mm in diameter with a length of 51 mm and was encased in a thin-walled aluminum housing with internal PTFE reflector; a schematic of the preliminary design concept for this

[1]These detectors were graciously provided for the purpose of this measurement by Prof. John Mattingly of the North Carolina State University Nuclear Engineering Department.

Detector number	Scattering angle (degrees)	Standoff distance (cm)	Rel. height (cm)
6	19.9	93.2	1.1
4	25.1	92.1	−0.5
9	31.3	92.4	1.2
5	37.4	90.5	0.4
10	44.0	92.3	0.6
7	55.2	78.0	−0.2
11	55.8	89.8	0.0
1	68.0	75.1	0.0
0	73.8	69.2	0.4
3	85.5	64.6	0.4
2	90.0	67.3	0.8
8	97.1	62.0	0.2

Table 4.1: Measurements of the backing detector positions relative to the scattering detector for the CsI[Na] quenching factor measurements. Relative heights reflect the vertical offset between the midpoints of the CsI[Na] scatterer and a given backing detector. Scattering angles assume the incident and scattered neutrons are in the same vertical plane, neglecting height differences between the backing detectors and the CsI[Na] crystal. Uncertainties on standoff distances and relative heights are all ±0.1 cm; angular uncertainties are ±1.9°, dominated by the size of the backing detectors.

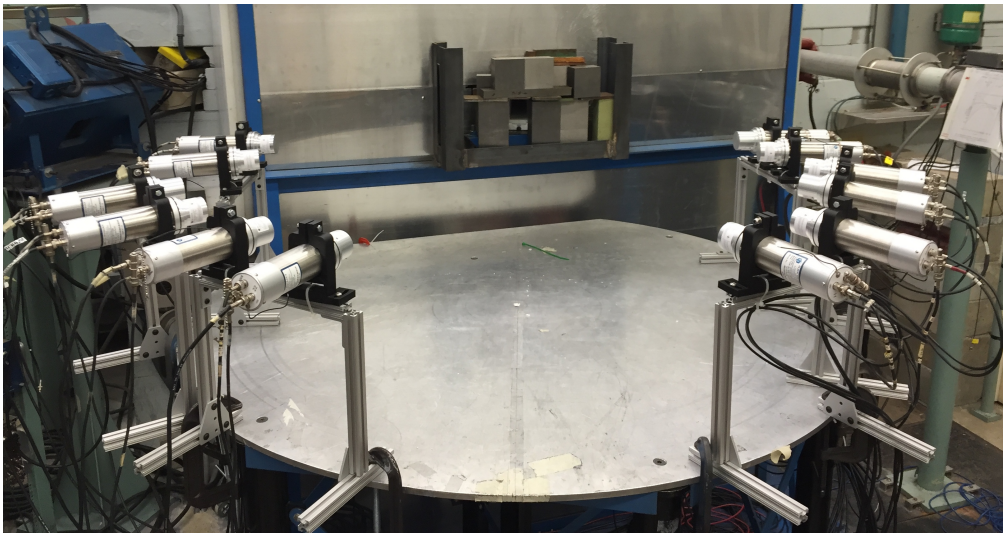

Figure 4.5: Photograph of the CsI[Na] quenching factor experiment setup in the TUNL SSA. The backing detectors, 2"-by-2" EJ-309 liquid scintillator cells, can be seen oriented towards a central scattering-target location; the scattering angles and standoff distances of the 12 backing detectors are shown in Table 4.1.

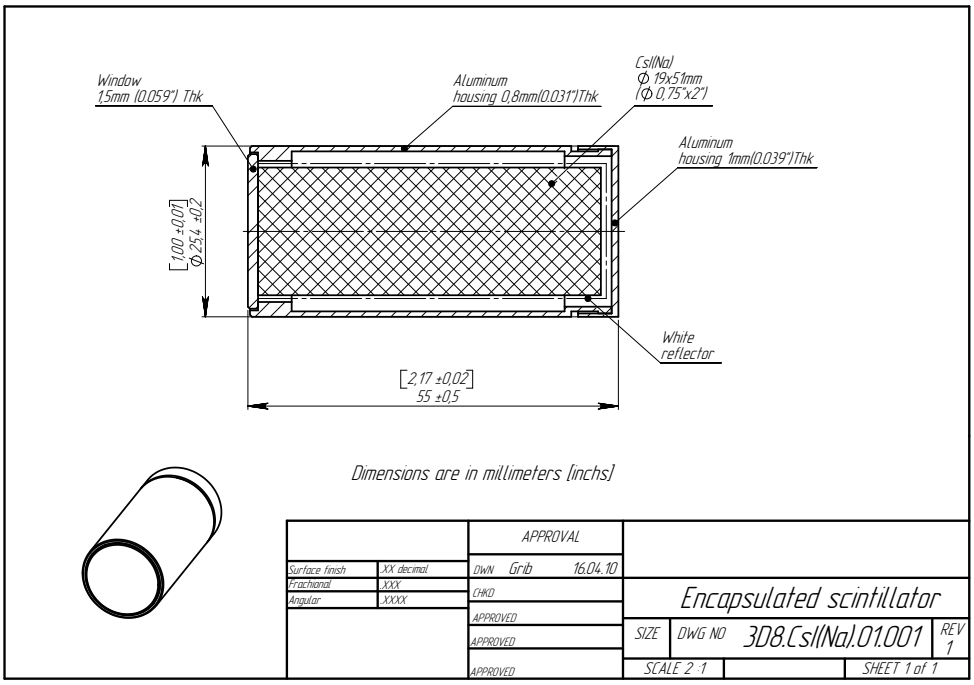

Figure 4.6: Schematic of the preliminary design concept for the small CsI[Na] detector assembly used for the QF measurements described here. Schematic provided by Philip Parkhurst of Proteus, Inc. [142].

detector assembly can be seen in Fig. 4.6. Through the center axial region of the assembly, the aluminum casing had a thickness of 0.8 mm with a PTFE thickness of 2.4 mm. The CsI[Na] scattering detector was mounted to a small, square-shaped PMT with an ultra-bialkali photocathode, which manufacturer Hamamatsu reports has a quantum efficiency of $\sim$40% for scintillators whose emissions are comparable in wavelength to those of CsI[Na] [208]. A 3-D printed mounting jig made of acrylonitrile butadiene styrene (ABS) was mated to a goniometric platform and supported the assembly of the PMT and CsI[Na] detector; Figure 4.7 shows an image of this system in place at the SSA. The detector depicted in Fig. 4.7 shows some deviation from the schematic shown in Fig. 4.6; appropriate dimensions for simulation are inferred from measurements of the assembly and comparison with schematics for similar assemblies which reflect the realized top-cap configuration.

4.2.2: Electronics and data acquisition

The anode output of the PMT mounted to the CsI[Na] scattering detector was sent through a low-loss coaxial cable into a Philips Scientific 771 [228] fast linear amplifier set to a gain factor of 10. Phototube gain was relatively high, as the PMT was operated at a bias voltage of -950 V, near the maximum recommended bias of -1 kV, but the amplifier was enlisted to further separate single-photoelectron signals from pedestal.

Figure 4.7: Photograph of the CsI[Na] detector assembly and its associated PMT mounted for the QF measurement at the SSA. Comparison of the realized assembly and its concept design, shown in Figure 4.6, shows some deviation from the original design. The end of the SSA neutron-beam collimator is visible in the background.

PMTs should be an exceptionally clean source of amplification and it is possible that preliminary examination of the data, without fully-developed algorithms to deal with baseline fluctuations, gave a false impression of insufficiency from phototube gain alone. The output of the PS771 amplifier was sent directly into a channel of the digitizer.

Timing of events relative to the beam pulse was enabled through the digitization of the bipolar signal from the beam-pickoff monitor (BPM) system. Physically, the BPM consists simply of a thin, cylindrical-shell of copper which is mounted concentrically in the charged-particle beam line. As the charged-particle beam passes through the $\sim 1^1/_2$"-diameter cylinder, a bipolar current pulse is induced and amplified just outside of the beamline by two cascaded, wideband Phillips Scientific model 6954 GHz bipolar amplifiers [227]. The amplifier output is tee'd, and one copy is sent via low-loss coaxial cable to the tandem accelerator control room while another is sent directly into a digitizer channel. In the control room, the BPM signal is sent through a TUNL-standard timing circuit which is functionally based on a zero-crossing discriminator; the output of this circuit is used to tune the accelerator system.

Signals from the anodes of the backing detectors were sent into individual channels of Mesytec MPD-4 [183] modules. The MPD-4 performs discrimination, amplification, and pulse-shape discrimination, produc-

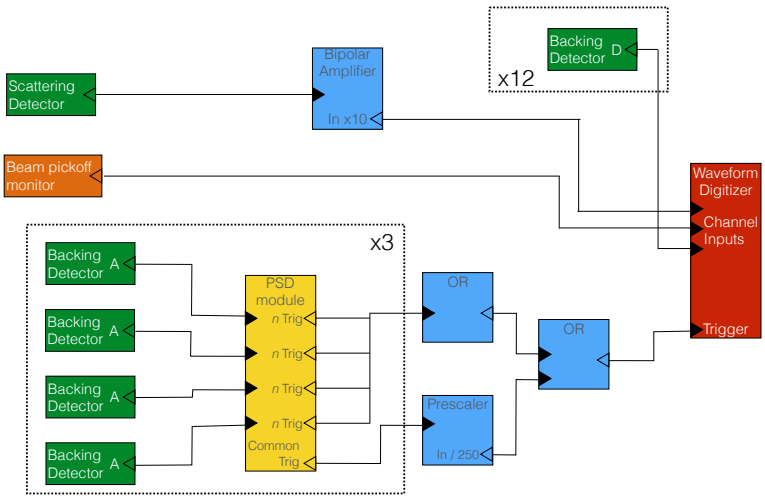

Figure 4.8: Diagram of data-acquisition electronics for the CsI[Na] quenching factor measurement.

ing analog voltage outputs proportional to the anode charge integral and the pulse-shape parameter of the pulse; a logical trigger pulse is produced on an individual-channel basis and a module-wide trigger, which can be configured to identify events of particular PSP character, is also available (see Sec. A.2.2 for additional discussion of the MPD-4 module). A total of four MPD-4 modules were in the system to cover 12 backing detectors and a single beam monitor (§4.3.1). Dynode outputs of the backing detectors were tee'd and locally terminated through 50 Ω with the parallel tee output going directly into a channel of the digitizer, with each backing detector having a dedicated channel.

Data was collected in the form of digitized waveforms using a CAEN V1730, a 14-bit VME digitizer sampling at 500-MS/s [67]. In total, 14 channels were digitized, representing the scatterer (CsI[Na]), the beam-pulse monitor (BPM), and the dynode signals from each of the 12 backing detectors. Each waveform was 30 μs long, configured such that the trigger occurred at $\sim +5.5$ μs with respect to the start of the waveform. Readout of the digitizer was accomplished via a PCIe-VME bridge interface card connected to the digitizer via fiber-optic link.

To create a trigger for the digitizer, the triggers from the MPD-4 modules indicating an event with neutron-like PSD were ORed together, along with an OR of *all* triggers in the backing detectors prescaled by a factor of 1/250. This trigger then represents a "neutron event" trigger with some γ-like events intentionally kept to provide a background sample.

A simplified diagram of the electronics can be seen in Fig. 4.8.

Line	Element	Energy (keV)	Intensity
$K_{\alpha 1}$	Cs	30.973	100
	I	28.612	
$K_{\alpha 2}$	Cs	30.625	53–65
	I	28.317	
$L_{\alpha 1}$	Cs	4.289	~ 90
	I	3.938	
$L_{\beta 1}$	Cs	4.619	100
	I	4.221	

Table 4.2: Characteristic X-ray energies and absorption lines for cesium and iodine. Values from Reference [192].

Section 4.3: Calibrations and stability

4.3.1: Neutron beam energy measurement

The energy of the neutron beam was extracted from time-of-flight data, representing the separation in time between the BPM signal and the detection of a neutron in a dedicated monitor detector. This zero-degree detector (so-called for its location at 0° with respect to the neutron beam axis) was operated for the duration of the experiment and, during a dedicated effort towards beam energy characterization, was moved to 3 different standoff distances from the neutron production target.

Neutron detection and the determination of the neutron beam energy, carried out using a Markov chain Monte Carlo fitting routine, are discussed in detail in App. D. The neutron energy distribution determined in this analysis is shown in Fig. D.6.

4.3.2: Energy calibration and resolution

Energy calibration of the CsI[Na] scattering detector was established using an ^{241}Am source, which has a prominent γ-ray of 59.536 ± 0.001 keV [154] along with numerous other γ and X-ray lines at lower energies and intensities [71]. In addition to the full-energy peak associated with the 59.54-keV γ-ray, atomic effects of the Cs and I constituents give rise to addition features in the observed spectra [154]. Pertinent characteristic X-ray lines for cesium and iodine are collected in Tab. 4.2.

L-shell escape peaks for both isotopes are poorly resolved and are treated as a single "tail" feature in the spectrum. The analysis focused only on the full-energy peak region, neglecting features of the K shell. Gaussians representing the full-energy peak and the "tail" were added to a uniform background and the data was fit in the region between 720×10^3 and 1400×10^3 ADC units. The amplitude of the tail feature,

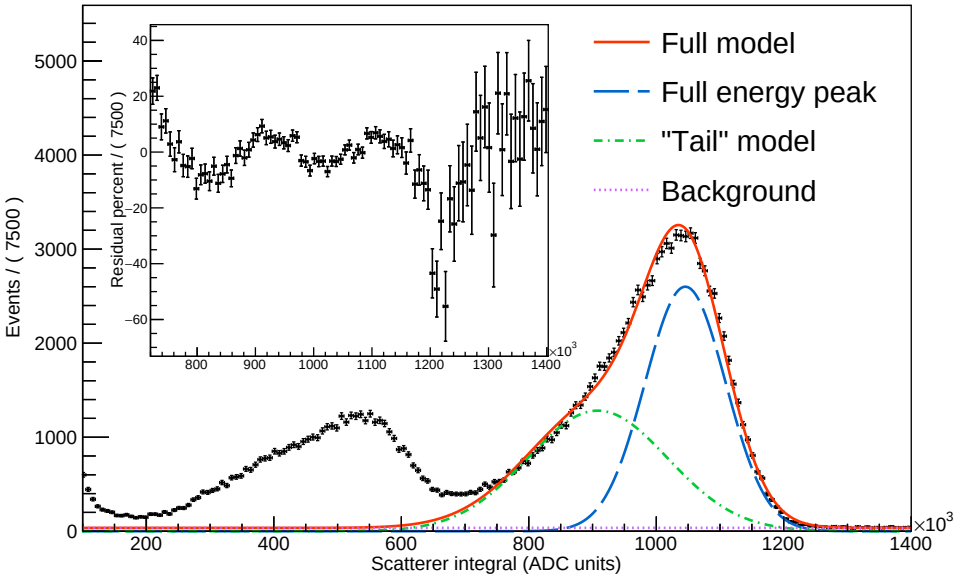

Figure 4.9: Spectrum collected with a ^{241}Am source in a close geometry with the CsI[Na] scattering detector. Several features are distinguishable: the structure from $\sim 800 \times 10^3$ – $\sim 1200 \times 10^3$ ADC units is composed of the full energy peak, L-shell escape lines, and contributions from downscattered full-energy γ rays; lower-energy structure represents K-shell peaks and lower-energy γ rays from the ^{241}Am source. The analysis required determination of the light yield only at the full-energy peak location and simplified the treatment of other spectral features in this area, neglecting features near the K-shell structure. The inset shows the residual of the data with the fitted model.

a very approximate representation of the combination of L-shell escape features and slightly downscattered 59.54-keV γ-rays, was not constrained in the fitting process. Figure 4.9 shows data and the fitted model for one of the calibration runs.

Light yield is treated as a linear function of deposited (electron-equivalent) energy defined by points at $E = 0$ and $E = 59.54$ keV. The assumption of linearity simplifies the presentation of the QF and allows normalization to be carried out relative to experimentally accessible references, should other experimentalists hope to adopt the results of the present study.

Source measurements with ^{241}Am were conducted at irregular intervals through the experiment, both immediately before and immediately following rotations of the CsI[Na] detector. Separately for each measurement, the fitting procedure described here was carried out. The determined calibration values showed good stability and agreement, and a single global calibration value was adopted for use at all times during the experiment. Section 4.3.4 discusses the determination of this global calibration value.

4.3.3: Calibration of single photoelectron charge

Proper analysis of the quenching factor data relies on the ability to determine the number of photoelectrons present in the recorded signals. Such a determination could be approached in numerous ways, including detailed analysis of individual events and the use of Bayesian statistical techniques to establish photoelectron counts and *times* of individual photoelectrons within large, many-PE pulses [12]. The analysis utilized here is comparatively simple and relies entirely on spectra of charge-pulse integrals. Conversion of these spectra into PE number can be accomplished by a straight-forward division by the mean charge of a single PE, assuming linearity in photoelectron charge across the analyzed range.

Independent values of the average single-photoelectron charge are determined for each hour-long acquisition period throughout the run. The CsI[Na] waveform for each event is filtered, as discussed in App. C, and then integrated over a 3.7-μs-long pretrace region. Pretrace integrals are histogrammed and the resulting spectrum for each run is fit using a physically-motivated model which additively combines the shapes associated with several low-charge features. Specifically, the model consists of a Gaussian which accounts for the pedestal (i.e., integration of "nothing"), an exponential which models various "noise" effects such as incomplete amplification through the PMT dynode chain [93], and the sum of the first few PE shapes.

The Gamma distribution is employed to model the charge distribution associated with single photoelectrons. PDFs for nPE shapes reflect the convolution of n SPE shapes, and thus the SPE PDF parameters entirely define the distributions for nPE signals. See App. B for discussion on the parameterization of single- and multiple-photoelectron shapes.

4.3.4: Stability of calibrations and determination of global calibration values

For each of the calibration parameters, values for each individual run can be determined. Variation of these values was relatively small over the course of the experiment. To reduce the analytical complexity introduced by individual-run calibration values, single values for the calibrations were adopted when aggregating the experimental data. Representative values for light yield at the ^{241}Am peak and the integral of SPE signals were found to be 29.88 ± 0.39 PE/keVee and 1502^{+22}_{-26} ADC/SPE, respectively, and the distribution of the values as a function of time are shown in Figs. 4.11 and 4.12. Subsequent discussion focuses on the method by which these global values are determined.

To determine a representative "global" value for calibration parameters, consideration of the philosophy behind such an approach is important. The fundamental question is: what single value and error for, or (similarly) distribution in, a given parameter accurately describes the individual observed values taken in aggregate? A global value or distribution serving this end will necessarily have larger errors, or be wider,

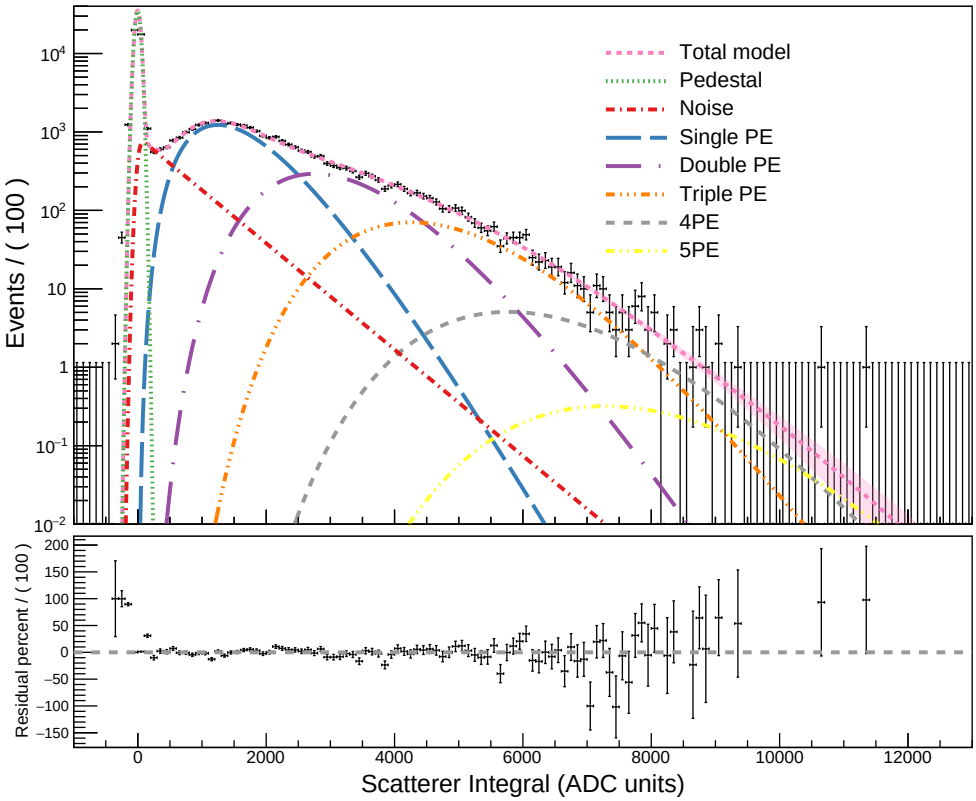

Figure 4.10: The distribution of integrals of the CsI[Na] waveform over the signal region, shown in ADC units, for a single hour-long run zoomed in to the few-photoelectron level. The data is fit with a model consisting of an additive combination of: Gaussian pedestal; exponential noise convolved with the pedestal; and Gamma-distributed models of nPE populations for $n = 1, \ldots, 5$. A shaded region around the total model, shown in hot pink, indicates the 1σ uncertainty band of the fit. The bottom panel shows the residual distribution of the fit, in percent of each data point, with a dashed gray line drawn at 0 residual. The model does not accurately fit a low-integral feature of the pedestal but describes the data well over the broad region with ample statistics. For a more complete description of the fitting procedure see Section 4.3.3 of the text.

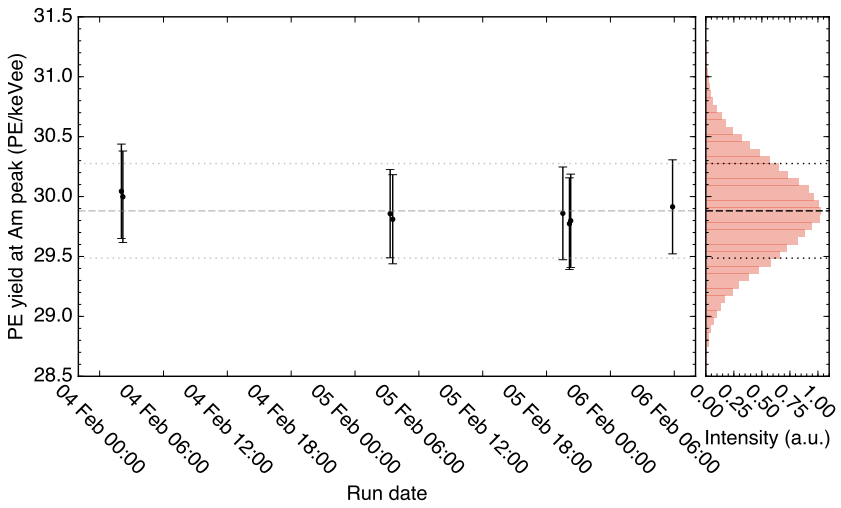

Figure 4.11: Measured "yield" for the CsI[Na] detector used in the QF measurements, in integrated digitizer units, for the full-energy peak of ^{241}Am, $E = 59.54$ keV. Source measurements were taken at several times through the experiment. The panel at right shows a kernel density estimate of the "global" value for the calibration parameter, determined using the method described in 4.3.4, finding a value of 29.88 ± 0.39 PE/keVee.

than the errors on individual data points.

Temptation to perform a *fit* through the measured parameter values should be avoided. Though the result of a fit of a constant value through a series of parameter data will indeed yield a single value with (speciously) appropriate error and quite probably acceptable results from any chosen goodness of fit tests, *the underlying question that is answered by a fit is different and not compatible with the goal of determining a global, representative parameter value.* Interpreted through the more-intuitive Bayesian view of probabilities and statistics, a fit responds to a scenario posed thusly: a single, "true" value of the parameter exists and the data are the result of different measurements of this single value; given the data, what is the underlying value of the parameter? With this philosophical underpinning, it follows naturally that the fitted value of a parameter will generally have smaller error than the individual data points[2].

To establish global parameter values which more accurately reflect the real variation of calibration values reflected in the observed data, an approach similar to a Gaussian kernel density estimate is used [231]. For each individual calibration parameter data point β_i with error δ_i, where $i = 1, \ldots, N$ with N the number of calibration values measured, 1000 samples are drawn from a Gaussian PDF with $\mu = \beta_i$ and $\sigma = \delta_i$. The collection of N*1000 values are then taken to provide a non-parametric model of the desired "global" distribution of the parameter.

To utilize these global calibration parameter values in the analysis chain, we seek simplified representa-

[2]In the context of distributions or Bayesian approaches, the result of a fit will yield a more narrow posterior distribution on the parameter compared to the individual measurements.

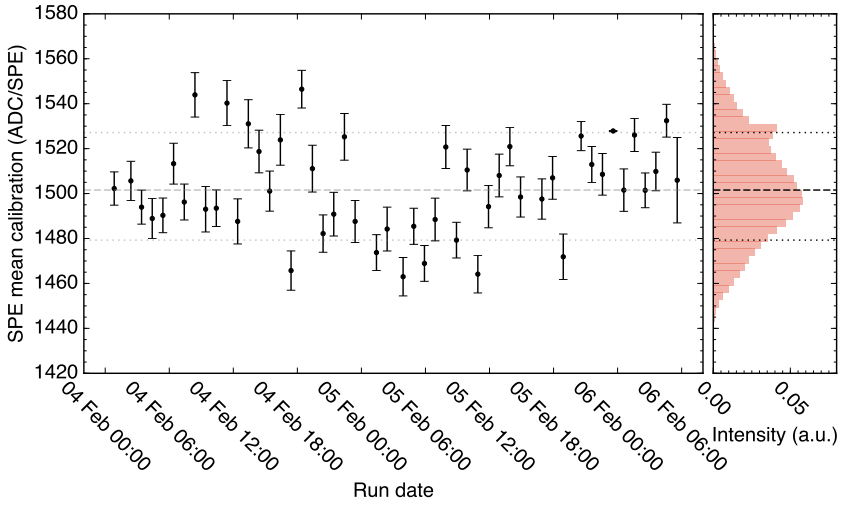

Figure 4.12: Average integrated charge (ADC units) per single photoelectron (SPE) for each hour-long run in the CsI[Na] QF experiment. The data analyzed to inform these calibrations is drawn from the integrals of the signal regions in waveforms from each individual run. The distribution of the "global" calibration parameter, found to be 1502^{+22}_{-26} ADC/SPE, is shown at right with the mean and a band corresponding to the central 68% of the distribution indicated. For a full discussion of the procedure for determining the SPE charge, see Section 4.3.3 of the text; for a discussion of the production of the global distribution, see Section 4.3.4.

tions of these distributions in the form of a single value with asymmetric errors,

$$\beta = \beta_{\mu}{}^{+\delta_+}_{-\delta_-}. \tag{4.3}$$

We take β_μ to be given by the mean of the distribution, δ_- to be the value below which 16% of the distribution falls, and δ_+ to be the value above which 16% of the distribution falls[3]. The distributions of "global" parameters, along with central values and central 68% confidence intervals, for the energy calibration and mean charge of SPEs can be seen on the right panels of Figures 4.11 and 4.12, respectively.

Note that the SPE calibration value used in subsequent analysis is 1486^{+20}_{-18} ADC/SPE, corresponding to a value derived from an earlier, less-constrained fit to the SPE shape. This is a ∼1% difference from the value derived using the model like that shown in Fig. 4.10, and the calibration values agree within uncertainty.

[3]In a fully Bayesian analysis chain, the aggregate numerical distributions resulting from Gaussian sampling of each data point could be utilized directly, without reduction to a single, representative value and uncertainty band.

Section 4.4: Simulations

4.4.1: MCNPX-PoliMi simulations of neutron beam measurements

Simulations of a geometry closely approximating the experimental configuration were carried out using MCNPX-PoliMi [204]. In the CsI[Na] scatterer itself, the sodium dopant is omitted as its low concentration results in negligibly few scattering centers compared to the Cs and I constituencies. A PTFE reflector surrounds the CsI volume and the assembly is completed by an external aluminum shell; these are modeled with thicknesses of 1.4 mm and 0.8 mm, respectively. Backing detectors are represented simply by volumes of the EJ-309 liquid scintillator, all positioned vertically such that the center points of the detectors and the CsI scatterer are matched to the height of the neutron beam.

Material definitions and properties were taken, in many instances, from Ref. [180]. The neutrons are produced according to a spectrum derived from the TOF fitting procedure outlined in App. D, with the spectrum of particular discussion in Sec. D.7.2; the utilized neutron spectrum is shown in Fig. D.6. The output of the MCNPX-PoliMi simulation was converted into ROOT trees for convenient analysis.

4.4.2: Backing detector events

A faithful modeling of the experiment requires consideration of the light-yield properties of EJ-309, the liquid scintillator used in the backing detectors. Signal generation in the more general class of organic scintillators is discussed in Sec. A.1, but the specific, experimentally observed characteristics of EJ-309 are the subject of Sec. A.1.3. In analyzing the simulations, the "observed" signal for proton recoils in the EJ-309 cells is quenched according to Eq.(A.1); this quenching is applied on an interaction-by-interaction basis, as multiple scatters in the scintillator cannot be treated as a single, aggregated energy deposition due to the nonlinearity of the light yield. Carbon recoils are quenched using the same parameterization but with an additional quenching factor of 20%.

Following the determination of the electron-equivalent energy for a given deposition, finite resolution effects are applied using a functional description of the energy-dependent resolution from Ref. [223],

$$\frac{\Delta E}{E} = \sqrt{\alpha^2 + \frac{\beta^2}{E} + \left(\frac{\gamma}{E}\right)^2},$$

with parameter values from Enqvist *et al.*, $\alpha = 0.113 \pm 0.007$, $\beta = 0.065 \pm 0.011$, $\gamma = 0.060 \pm 0.005$ [103]. These resolution parameters were determined for a 7.6-cm by 7.6-cm cylindrical detector and may not reflect the realizable resolution for the smaller detectors used here [103]. The authors of Ref. [103] do find that detectors of different sizes have meaningfully distinct resolution parameters, but only two detector sizes were

compared (12.7-cm by 12.7-cm cylinders and the aforementioned 7.6-cm cylinder) and at relevant energies the relative difference between the determined resolution functions was only $\sim 10\%$. In the case of the simulated data under discussion, the most "relevant energies" are ~ 1.5 MeVee, the endpoint energies of neutron events in the backing detectors. This region provides information for determination of the backing detector gain.

Cuts placed on the experimental data, discussed in Sec. 4.5.4, select a region of signal space where the neutron and γ-ray event populations are highly distinguishable. Consequently, the PSD performance of the backing detectors is omitted from consideration when processing the simulated interactions in the EJ-309 cells.

Section 4.5: Analysis approach for digitized data

Digitized data possesses, for each event, a scatterer waveform, a BPM waveform, and waveforms for the dynodes of each of the backing detectors. Analysis of this data makes use of information in individual waveforms as well as timing information between channels for a given event. The outline for analysis of each event is as follows:

1. Process each of the 12 backing detector waveforms, performing simple pulse finding.

2. Refine timing of each detected backing detector pulse; carry out integration, and PSP determination for each.

3. Determine the backing detector "trigger pulse" - which one (in the case of more than one pulse) triggered the digitizer for the event.

4. Determine the BD-BPM timing: starting at the time where the trigger pulse occurred, step forward in the BPM waveform until the value exceeds a nominal ADC value above "baseline"; then, look for the waveform value to fall below baseline, indicating a zero crossing has occurred; determine the time of BPM zero crossing, and take the difference between the BD pulse time and the BPM zero crossing time.

5. Starting from a time defined with respect to the backing detector trigger time, begin searching the CsI[Na] waveform for the first photoelectron signal[4].

6. Integrate the CsI[Na] signal for 3 μs, starting at the first SPE time.

[4] "First" refers to the first photoelectron signal in the search region. This will not necessarily be the first SPE observed in the waveform.

This is somewhat reductive but serves to present in "broad strokes" the general approach. Additional details of each stage and at an event level are discussed in the following sections.

4.5.1: Treatment of backing detector waveforms

Each digitized event contains waveforms from the dynodes of all 12 backing detectors; prior to offline analysis, the backing detector associated with triggering the event is not known and each of the 12 waveforms are processed to make such a determination. Baselines and estimates for the FWHM of the baseline noise are established for each backing detector using the "normal mode approximation" discussed in Sec. C.4.2; a large fraction of the waveform is used in this determination, starting 10 samples into the waveform and including the next 14,000 samples (the entire waveform is 15,000 samples long). Pulses are identified in each backing detector using a edge-detection algorithm which looks for instances where the waveform exceeds the baseline estimate by at least 7× the estimate of the standard deviation of the baseline fluctuations. For each detected pulse in any of the backing detectors, a CFD algorithm is used to determine the time at which the pulse passes through 20% of the maximum pulse amplitude. This timing information is used in any subsequent timing-related analysis.

The integral of each pulse is recorded over a period 206 ns in duration, beginning 6 ns in advance of the CFD time. Pulse-shape discrimination was carried out using the charge-integration method (§A.2.1) and a pulse-shape parameter value like that described in Eq.(A.2) was established for each event. In the determination of the PSP, the "full" region was the same used in determination of the pulse integral and the "tail" region began 16 ns after the CFD time, sharing the same end boundary as the full region (200 ns following CFD). When calculating amplitude, integral, and PSP, the normal mode approximation of the baseline is utilized. Figure 4.13a shows a representative example waveform from one of the backing detectors, the baseline estimate, and the relevant timing features. Data from one of the backing detectors collected during the QF experiment and analyzed using the described integration windows is shown in Fig. 4.13b; though there are some variations in the relative intensities of γ- vs. neutron-like events between backing detectors, the general characteristics of the data shown here is representative of that from all detectors.

Information on all of the pulses detected for each event is calculated and preserved during the initial analysis phase; this information includes the identity of the backing detector in which the pulse occurred, timing, amplitude, integral, and PSD.

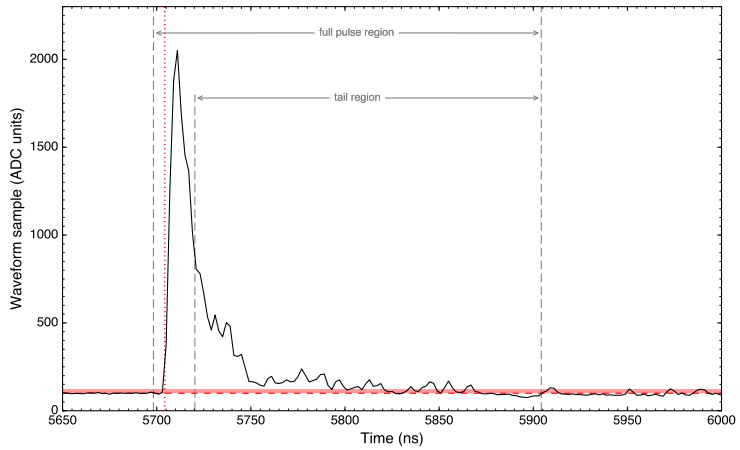

(a) Example waveform for backing detector 6 with PSP-relevant integration regions indicated

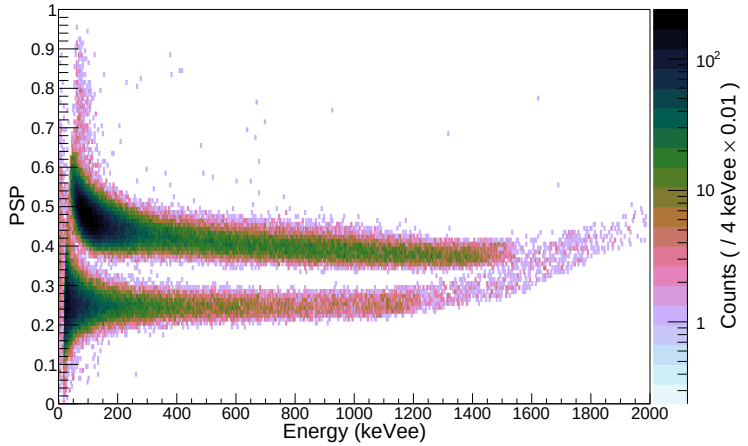

(b) PSD plot for backing detector 6

Figure 4.13: Example of backing detector waveforms and PSD performance. Panel (a) shows a representative backing detector pulse. The red, dashed horizontal line indicates the baseline estimate from the normal mode approximation and a lightly shaded red region above the baseline indicates the region below the peak-detection threshold used during the analysis, set to a value of 7× the estimated standard deviation of the baseline fluctuations. The red, dotted vertical line indicates the pulse time as determined by the CFD algorithm and gray, dashed vertical lines at successively larger time values represent the start of the full pulse region, the start of the tail region, and the end of the pulse region, respectively. The ratio between the integrals of the tail region and the full pulse region define the pulse shape parameter used to discriminate between neutron- and γ-ray-like events in a parameter space like that shown in (b), which plots the PSP against full pulse integral showing two distinct populations: neutrons are centered at PSP ~ 0.45 while γ events occur with lower PSP values, ~ 0.25. Data shown here are from backing detector 6, but reflect the general behavior of all backing detectors.

4.5.2: Treatment of beam-pickoff monitor waveforms

Waveforms from the beam-pickoff monitor (BPM) system were digitized and allowed beam-pulse relative timing of events in both the CsI[Na] and backing detectors to be determined. Signals from the BPM system represented current induced as the charged-particle beam passed axially through a conducting cylinder (§4.2.2) and the zero-crossing time is taken to represent the "beam time". As can be seen in Fig. 4.15, each digitized event contains many BPM pulses. The timing for an event does not need to be established relative to "the" BPM pulse which actually was produced by the same beam bunch which participated in the event; because of the periodicity of the BPM pulses, relative timing to the first BPM pulse following the backing-detector trigger pulse is sufficient to unambiguously describe events. Beginning to search at the trigger-pulse time, the algorithm would seek the sample at which the BPM waveform crosses under its baseline value after having exceeded baseline by 500 ADC units. The midpoint time between the two samples straddling the baseline value is taken as the zero crossing time.

4.5.3: Properties and treatment of CsI[Na] waveforms

Consideration of the nuclear-recoil energies under exploration, the quenching factor for nuclear recoils from earlier measurements, the light yield of CsI[Na] ($\sim$ 41 photons/keVee [81]) , and the approximate quantum efficiency of the PMT ($\sim$ 40% [208]) leads to a recognition that the events in the scattering detector should consist of between $\sim$ 5 and $\sim$ 85 photoelectrons, depending on the backing detector into which the beam neutrons are scattered. As CsI[Na] is a moderately slow scintillator, with a single-exponential approximation of its timing characteristics having a decay time of $\sim$ 675 ns [77], it becomes evident that "events" in the CsI[Na] detector may be manifest as a series of entirely distinct single-photoelectron pulses in the waveforms, especially at the lower scattering angles. Integration of the waveform then consists of integrating long periods of signal-free baseline, which can make the resulting spectra sensitive to fluctuations in the baseline and the way these are addressed in the integration routine. Appendix C presents a discussion on this subject and outlines the conditional moving-average (CMA) filter which is used subsequently to mitigate the impact of baseline fluctuations on spectral resolution. For the analysis of the QF data presented here, the half-width of the CMA filter was set to 100 samples (200 ns). A baseline value along with an estimate of the FWHM of fluctuations about this value was produced using the normal mode approximation (§C.4.2); the CMA preload value was this baseline estimate, and the rejection threshold was $4 \times \mathrm{FWHM}/2.355 \approx 4\sigma$, using the estimated FWHM value from the normal mode approximation.

In addition to acquisition-related issues, the CsI[Na] waveforms can be contaminated by "afterglow" or afterpulsing from large-energy preceding events. This effect is well documented in alkali halide scintillators,

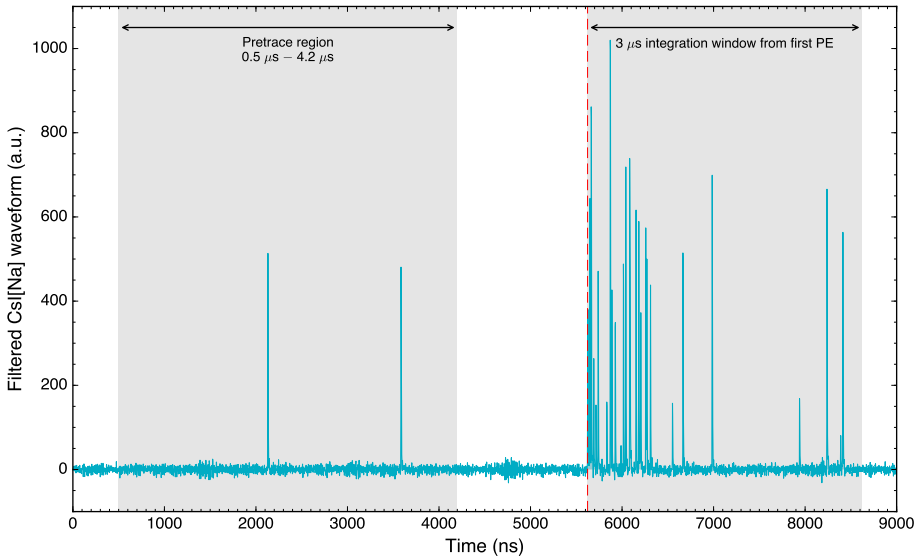

Figure 4.14: Detail view of the CsI[Na] scattering detector waveform shown also in Figures 4.15 and 4.16. Two integration regions are shown shaded in gray. The pretrace region is between 500 ns and 4200 ns of the waveform; this is used to identify events whose signal region may demonstrate contamination by long-lived light production associated with preceding energy deposition. The pretrace region is also used to populate a data set used for single photoelectron charge calibration. The signal region has a duration of 3 μs and begins at the CFD time associated with the first detected photoelectron in an appropriate arrival window. SPE arrival windows are defined relative to the time of the pulse observed in the triggering backing detector; see Section 4.5.4 of the text for a discussion of this timing.

and with decay times in the millisecond-range possible, an interaction which takes place well before the start of a waveform could still contribute significantly to the integrated signal attributed to an elastic scattering event [154]. To address this issue, a "prepulse" region is defined at early times in the digitized waveform from the CsI[Na] scatterer, extending from 500 ns into the waveform until 4200 ns. Both the integral of this region and the number of pulses found using a leading-edge detection algorithm are compared against threshold values: if more than a single pulse is detected in the prepulse region, the entire event is rejected from analysis. The integral of the prepulse region is used as additional protection against the acceptance of large amplitude events, which may appear to only have a single "edge" as the signal passes through threshold and remains above it, resulting in a large total integral; if the prepulse integral exceeds the equivalent charge of $\sim$ 26 PEs, the event is rejected.

The integration windows are highlighted over a CMA-filtered waveform from the CsI[Na] detector in Fig. 4.14.

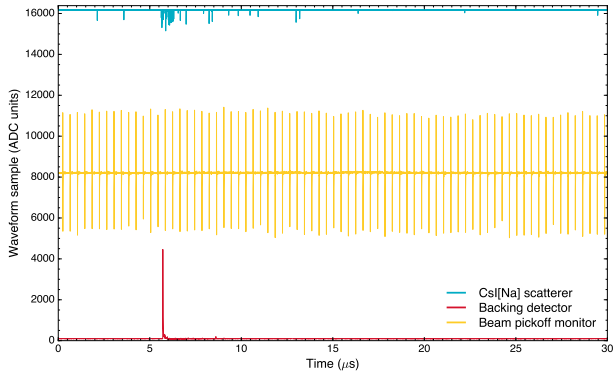

Figure 4.15: Waveforms from the CsI[Na] scattering detector, backing detector 7, and the beam pickoff monitor for a candidate elastically scattered neutron event. Digitized dynode signals from the other 11 backing detectors are omitted. This is intended to provide a qualitative impression of the event structure; detail of the timing region of the same event can be seen in Figure 4.16.

4.5.4: Full event consideration and selection of signal and background populations

Interpretation of any single event recorded during the QF experiment was based on the properties of the individual detector signals (i.e., scatterer, backing detector) as well as certain relative timing values between different channels. The pulsed nature of the neutron beam and the full digitization of signals from the beam pickoff monitor, the backing detectors, and the scattering detector itself result in an information-rich dataset with many possible ways of defining signal and background regions. In the end, requirements were based on the following properties:

- Timing between the backing detector CFD time and the following beam pickoff monitor zero-crossing time, referred to as "BD-BPM separation".
- Timing between the backing detector CFD time and the arrival time of the first photoelectron in the signal region of the CsI[Na] scatterer: "BD-SPE separation".
- Backing detector signal integral.
- Backing detector pulse-shape parameter value (PSD).

Figure 4.15 shows a full view of waveforms from an example event, including the scatterer signal, the triggering backing detector signal, and the BPM. Figure 4.16 uses the same set of waveforms, after applying the CMA filter to the CsI[Na], and illustrates the relative-timing parameters which were developed.

The BD-BPM separation metric in particular offers an attractive lens through which to view the data as the definition of a signal ROI relies only on the pulsed nature of the beam, the relatively narrow distribution of beam-neutron energies, and simple kinematics. Looking at the behavior of the PSP and integral distributions

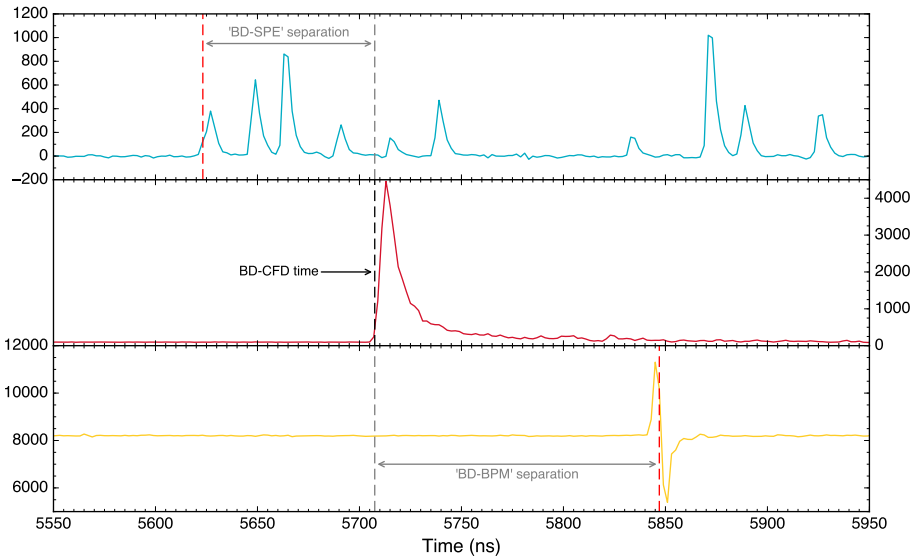

Figure 4.16: Detail view of timing region of the event also seen in Figure 4.15. The timing-related parameters used for event selection are overlaid.

of backing-detector events as a function of BD-BPM separation, a clearer picture of the characteristics of beam-neutron events can be established. Figure 4.17 shows a corner plot of the BD-BPM separation, the backing-detector pulse integral, and backing-detector PSD; the region of this parameter space associated with the pulsed neutron beam is clearly apparent.

The pulse-shape parameter distribution shown in Fig. 4.17 and the center bottom panel of Fig. 4.17 aggregate data over all BD-BPM times and already shows two distinct populations, in agreement with the expectation informed by the fact that neutron-generated pulses are more likely to have a larger fraction of "delayed" light (§A.1.2). Isolation of a beam-related region of time showing significantly heightened high-PSD-event intensity is in line with this understanding (see the bottom, right panel of Fig. 4.17). The total light observed in the backing detectors due to a beam-neutron event also shows a distinct feature in the larger fraction of higher-energy ($>\sim$ 300 keVee) events; see the center, left panel of Fig. 4.17.

To better explore the distinct characteristics of beam-related signals, we define two regions of time in BD-BPM separation space: a "beam-on" region and an "off-time" region whose boundaries are shifted by +100 ns with respect to those of the beam-on region. The beam-on region is defined on an individual backing-detector basis. Events in the off-time region are expected to be uncorrelated with the beam and represent "accidental" backgrounds that should be uniformly distributed in time, or equivalently, uniformly distributed through BD-BPM space with a statistically equivalent number of events contained in the two equal-length regions. Figure 4.18 shows the observed energies from these two regions for backing detector 7. From Fig. 4.18, it is clear that the accidental events, those present in the beam-off population, are largely

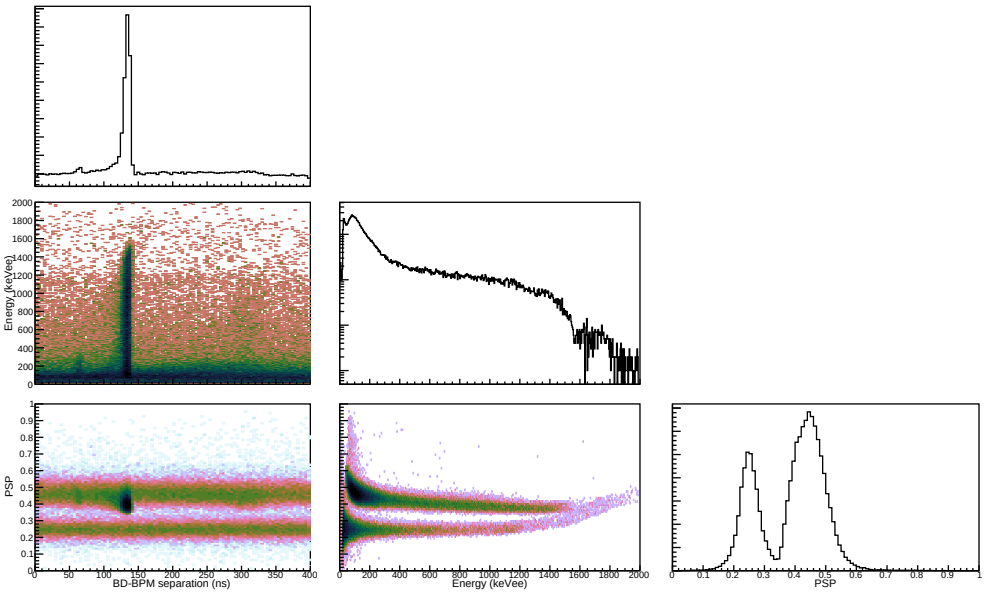

Figure 4.17: A corner plot showing characteristics of the events in backing detectors and their correlations. This data is for backing detector 6 but represents general properties of all backing detectors. Clear relationships exist between all parameters, and these can be utilized to select "signal" events.

confined to low energies. Introducing a requirement that the observed energy in the backing detector exceed a certain level should then preferentially select beam-related events and reject uncorrelated backgrounds, providing a more pure beam-related data sample for analysis. On an individual backing-detector basis, integrated signal thresholds are chosen based on populations like those shown in Fig. 4.18; these thresholds ranged from $\sim$ 220 keVee to $\sim$ 400 keVee.

It is necessary to develop an understanding of any bias introduced through the addition of the backing-detector energy cut. To this end, the simulations from Sec. 4.4 are used to explore the dependence of the nuclear-recoil energy in the CsI[Na] on the observed energy in the backing detectors. Cuts at backing-detector energies of 200 keVee, 500 keVee, and 1 MeVee were applied to the simulated elastic-scattering distributions; the results of these cuts can be seen in Fig. 4.19 and show no appreciable bias on the recoil energy is introduced by the use of cuts on this parameter. Consequently, the uncertainty introduced by adoption of a cut on backing-detector signal integral is neglected.

A detector-specific PSP cut preferentially selecting neutron-like events was also developed. The effect of the backing-detector integral and PSP cuts on the BD-BPM data for a single backing detector are shown, both individually and combined, in Fig. 4.20, where they can be compared with an uncut distribution.

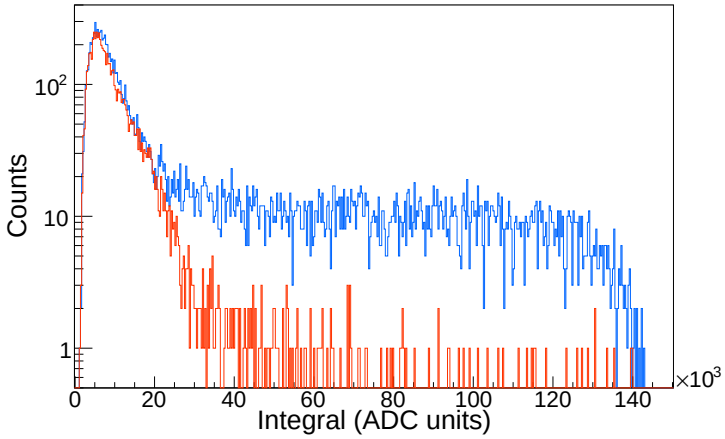

Figure 4.18: Distributions of the total integrated signal from backing detector 7 from two different regions of BD-BPM space: the beam region, shown in azure; and the off-time region, shown in red. The beam region shows a considerably higher fraction of events with larger integrated signals.

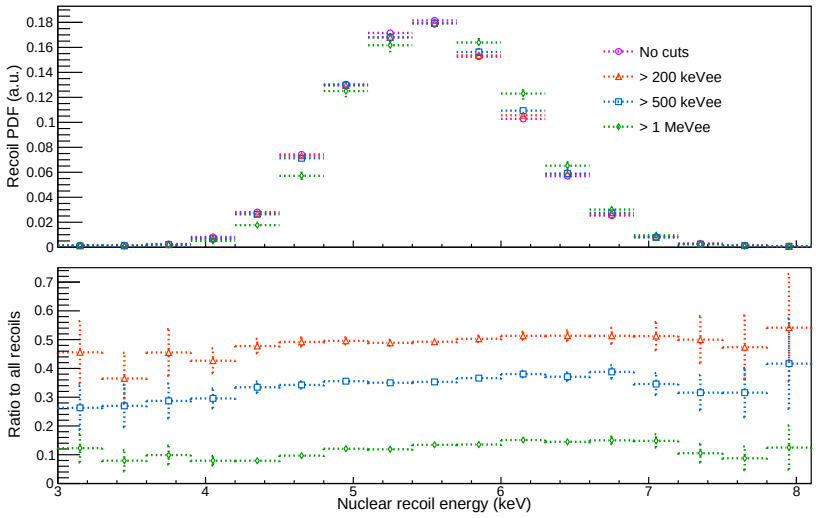

Figure 4.19: Simulated elastic-scattering recoil distributions associated with backing detector 6 with and without the application of backing-detector integral cuts at various electron-equivalent energies. The top panel shows the normalized distributions for all recoils and for those subject to cuts at 200 keVee, 500 keVee, and 1 MeVee total backing-detector energy; the distributions are equivalent within the statistical error shown except for the 1 MeVee cut which shows a small shift towards higher mean recoil energy, negligible compared to the width of the distributions. The bottom panel presents the fraction of total events selected by each cut; for each cut, the distributions show no significant deviation over the range of recoil energies. Both of these serve to demonstrate the minimal impact of recoil-energy bias introduced by use of the backing-detector integral cut.

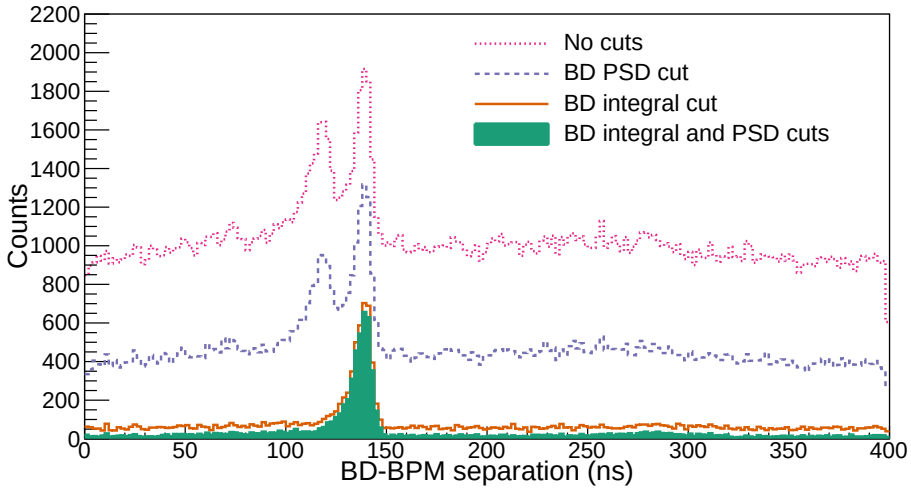

Figure 4.20: BD-BPM distribution for backing detector 7, showing the effect of cuts on backing-detector signal properties. The deep-pink, dotted line shows the distribution prior to application of any cuts on the features of the BD signal; requiring neutron-like PSP values yields the slate-blue, dashed line; the backing-detector integral requirement produces the dark-orange line; and the sea-green shaded histogram is produced by application of both the PSD and integral requirements.

Section 4.6: Fitting of spectra and extraction of quenching factors

4.6.1: Expected nuclear-recoil energy distributions from simulation

Using the simulations described in Sec. 4.4, distributions of the energy deposited in the CsI[Na] scattering detector by elastic scattering of beam neutrons were obtained for each of the 12 backing detectors. For each simulated event, the energy deposited in the CsI volume is recorded along with general properties of the interaction, including the incident particle and the kind of interaction (e.g., elastic scatter, inelastic scatter, capture); details about the scattered particle and its interaction with any backing detectors are also stored. Only those scatters associated with a (quenched) neutron energy deposition in the backing detector exceeding 200 keVee were kept, approximating the effect of the backing-detector energy cut (§4.5.4). An additional requirement that the backing-detector energy deposition occur within a detector-specific 30-ns window served to represent the cut in BD-BPM space. The resulting distributions can be seen in Fig. 4.21.

4.6.2: Distribution of expected observed signal yield

Translation of the expected recoil distributions of Sec. 4.6.1 into distributions in terms of observed signal requires consideration of non-trivial signal production and detection effects. Generically, these effects introduce terms contributing to decreased resolution of the scattering detector. Prior to incidence on the photocathode of the PMT, nonuniformities in the crystal and the coupling between the crystal and the PMT introduce a finite resolution that we refer to here as intrinsic. Intrinsic resolution effects are introduced

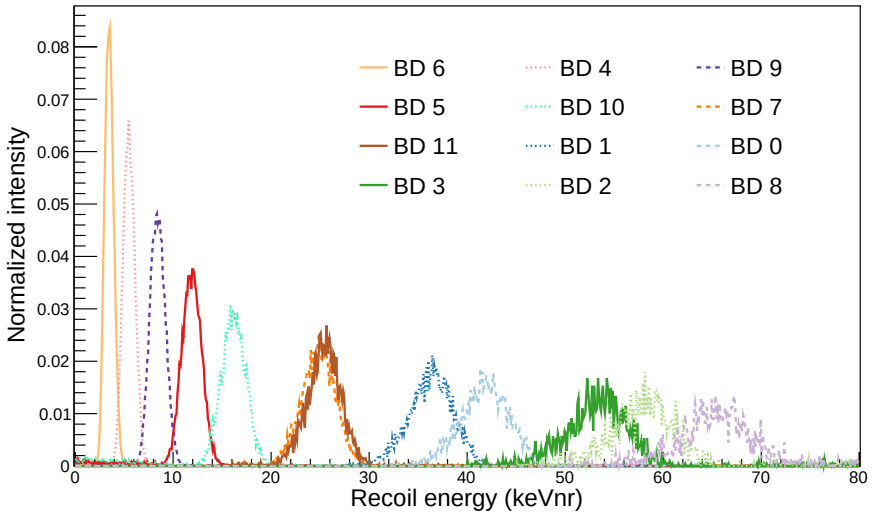

Figure 4.21: Expected nuclear-recoil energy distributions from the elastic scattering of beam neutrons by the CsI[Na] detector into the different backing detectors. These distributions are informed by MCNPX-PoliMi Monte Carlo simulations of a geometry closely approximating the experimental configuration, using an incident neutron-energy distribution following the experimental beam-energy measurement discussed in Appendix D.

as a Gaussian spreading of the nuclear recoil distribution with a width of $\sigma_{\text{int}} = \mu_{\text{NR}}^{\text{sim}} R_{\text{int}}$, where $\mu_{\text{NR}}^{\text{sim}}$ was determined earlier and R_{int} is allowed to float, with a starting value of 0.02.

Following the application of intrinsic resolution spreading, the model is in terms of nuclear recoil energies and is converted to the number of photoelectrons expected prior to consideration of statistical effects associated with PE production, referred to as the number of "raw" PEs. The number of raw PEs $n_{\text{PE}}^{\text{raw}}$ is found using

$$n_{\text{PE}}^{\text{raw}} = (QF_{\text{NR}}^{\mu}) \times Y_{\text{PE}} \times E_{\text{NR}}^{\text{sim}}, \tag{4.4}$$

where Y_{PE} is the photoelectron yield in units of PE / keVee and QF_{NR}^{μ} is the average nuclear-recoil quenching factor over the recoil energies represented by events in the relevant backing detector. Statistical effects are approximated by spreading the distribution in $n_{\text{PE}}^{\text{raw}}$ by a continuous approximation of a Poisson distribution with a mean of $n_{\text{PE}}^{\text{raw}}$, with the resulting distribution taken to represent the number of observed PEs n_{PE}.

4.6.3: Formulation of random-coincidence background model

An empirically determined background model was developed to represent the spectral component associated with events in which there is no observed signal from beam-related energy depositions in the CsI[Na]. Energy spectra from the CsI[Na] detector associated with events within the BD-BPM separation range of 340 ns – 400 ns whose PSP was γ-like were summed together to serve as the template for a background spectrum. Since it is assumed that this background is independent of which backing detector triggered the DAQ, data

from all backing detectors was collected to produce a relatively high-statistics template for the spectrum. In the process of aggregating events from all backing detectors, each individual event was subject to the PSP and observed-energy cuts appropriate for the backing detector which had triggered the system to record the respective event. As in the case of selecting signal data, events with more than a single participating backing detector were rejected.

The collected background data represented integrals expressed in ADC units. As the fit was ultimately carried out in the same unit space, no conversion or re-expression of the data (e.g., in units of photoelectrons) was necessary. To produce the PDF employed during the fitting procedure, the background data was histogrammed with uniform-width bins, though, due to considerable differences in the available statistics associated with any single backing detector, the bin-width differed depending on the recoil energy being analyzed[5]. Following appropriate normalization, the PDF was then described by linear interpolation between bin centers.

4.6.4: Complete spectral model and fit process

For each backing detector, the recoil model was additively combined with the background model and an extended maximum likelihood fit was performed on the signal-region data, allowing the counts in each component to float as well as the quenching factor. Uncertainties on the quenching factor fit value were determined using the `minos` algorithm [145] which determines 1σ error boundaries equivalent to the profile likelihood method, incorporating uncertainties associated with systematics and fit-results of other floating parameters (see §6.6.2 for a discussion on profile likelihood). The quenching factor entered the model in such a way that its variation would seem to "stretch" or "shrink" the recoil model in terms of observed photoelectrons; its best-fit value resulted in the most faithful reproduction of the observed photoelectron signal with the nuclear recoils predicted by the MCNPX-PoliMi simulation. Fitted models overlaid on data for several backing detectors can be seen in Fig. 4.22, including the most shallow (Fig. 4.22a) and most extreme (Fig. 4.22c) scattering angles. Results of the QF fits for each of the backing detectors are tabulated in Tab. 4.3.

Section 4.7: Quenching factor for a CEνNS search and comparison with literature data

Figure 4.23 shows the QF measurements from the present effort, earlier efforts in the literature (Refs. [126, 195]), and a measurement carried out by other members of the COHERENT Collaboration in close

[5]Since the background data was collected from all detectors, the *background* statistics do not change depending on which detector is being analyzed. The changes in binning are the consequence of the fitting machinery and compensation for variable statistics in the signal dataset.

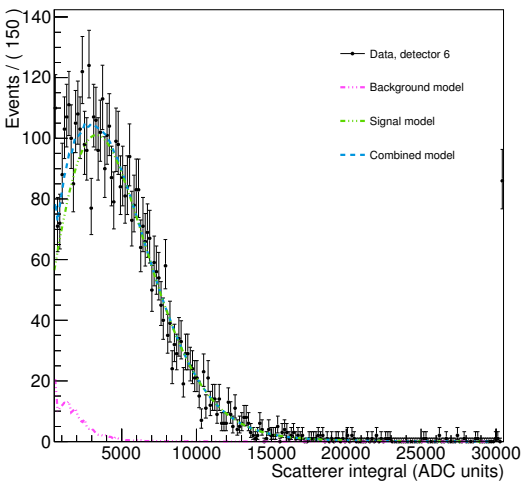

(a) Backing detector 6, $E_{\text{rec}} \approx 3.44$ keVnr.

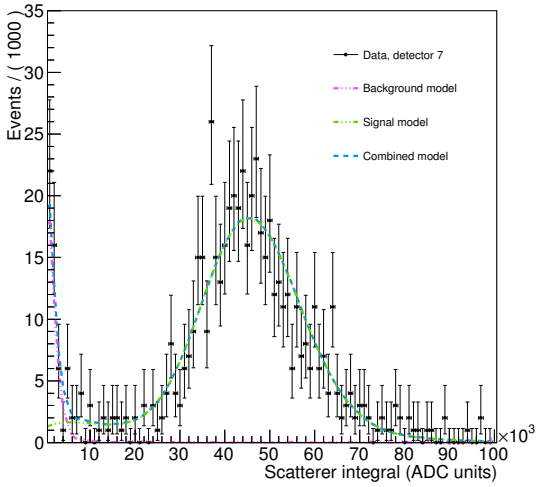

(b) Backing detector 7, $E_{\text{rec}} \approx 24.58$ keVnr.

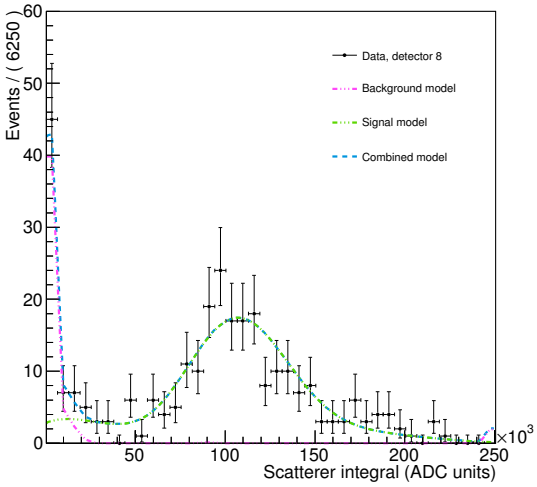

(c) Backing detector 8, $E_{\text{rec}} \approx 64.86$ keVnr.

Figure 4.22: Fitted models overlaid on experimental data for 3 different scattering angles in CsI[Na] QF experiment. Magenta lines, visible almost exclusively at low values of scatterer integral, represent the background model informed by out-of-beam-time data (see Section 4.6.3). The signal model, produced by applying response effects to Monte Carlo recoil-distribution spectra, is shown as a green line. An additive combination of the signal and background models, shown in blue, is fit to the data for each backing detector. Due to very different statistical content, the bin width used for different backing detectors varies.

Detector number	Recoil angle (degrees)	Recoil energy (keVnr)	QF (%)
6	19.9	$3.44^{+0.39}_{-0.36}$	$7.38^{+0.10}_{-0.14}$
4	25.1	$5.45^{+0.50}_{-0.52}$	$7.97^{+0.13}_{-0.12}$
9	31.3	$8.35^{+0.71}_{-0.69}$	9.32 ± 0.08
5	37.4	$11.78^{+1.00}_{-0.90}$	$10.10^{+0.09}_{-0.10}$
10	44.0	$16.0^{+1.42}_{-1.17}$	10.34 ± 0.13
7	55.2	$24.58^{+1.80}_{-1.65}$	11.00 ± 0.13
11	55.8	$25.17^{+1.86}_{-1.60}$	11.04 ± 0.18
1	68.0	$36.14^{+2.46}_{-2.07}$	10.76 ± 0.11
0	73.8	$41.60^{+2.93}_{-2.50}$	10.75 ± 0.12
3	85.5	$52.80^{+3.70}_{-3.20}$	10.40 ± 0.15
2	90.0	$57.58^{+4.72}_{-3.46}$	9.75 ± 0.20
8	97.1	$64.86^{+5.15}_{-4.23}$	9.67 ± 0.23

Table 4.3: Results of the quenching factor measurement of CsI[Na]. The recoil energy values are the means from the shapes seen in Figure 4.21; the uncertainties correspond to the central $\pm 1\sigma$ interval. Uncertainties on the recoil angles are all $\pm 1.9°$, with the diameter of the backing detectors the dominant contribution. Reported uncertainties on the QF values are determined by the `minos` algorithm [145] which incorporates uncertainties of other values in the fit model and represents the resulting 1σ error band.

collaboration with researchers at TUNL and using the same neutron beam facility as described here. All experiments normalize light yields in a similar way: using a two-point linear model between 0 and the light yield at ^{241}Am. The two COHERENT measurements of the QF are in agreement only at the lowest recoil energies measured; above ∼10 keVnr, the data series diverge. The Chicago data seems to flatten to a value of roughly 7%, while the TUNL measurement discussed in detail here continues to rise to a peak of ∼11% at 24 keVnr before declining slightly.

When the COHERENT measurements are considered alongside the existing literature values from Guo *et al.* [126] and Park *et al.* [195], little qualitative clarity is developed by the new measurements. The new data sets do show a trend towards decreasing QF at the lowest measured nuclear recoil energies, however; this is in contrast to the previous lowest measurements in Ref. [195], where a modest upward movement is observed. Models of the QF behavior, such as that by Tretyak [242], predict the upward trend suggested by the data of Ref. [195]. However, recent measurements of the QF for sodium-nucleus recoils in NaI[Tl] have shown decreasing values at low recoil energies like those observed here [75, 235, 257].

Subsequent analysis efforts will be dedicated to exploring the discrepancies between the two COHERENT data sets. The presence of a unaccounted-for systematic could explain the disagreement over much of the energy range. Resolution of this disagreement would provide higher confidence in the observation of downward-trending QFs at low recoil energies in inorganic, doped crystalline scintillators. Additional evidence for the failure of semi-empirical models based on Birks' law [242] would warrant additional exploration of the microphysics underlying the signal generation process. Some discussion of this can be found in Sec. 7.7.

To inform the analysis of the CsI[Na] CEνNS search data, a single, representative QF value was sought in the nuclear recoil energy ROI for CEνNS signals. The central representative value was determined by the weighted average of each data point which falls within the range of 5 – 30 keVnr while a conservative uncertainty estimate for the representative QF by taking the standard deviation of the same data points. This approach yielded a QF of $8.78 \pm 1.66\%$. Figure 4.23 shows the literature QF values along with the new measurements by the COHERENT groups; overlaid on this plot is a dashed, horizontal black line which shows the QF taken as representative of the global data over the ROI, with a shaded uncertainty band drawn in the region-energy range over which the value is relevant.

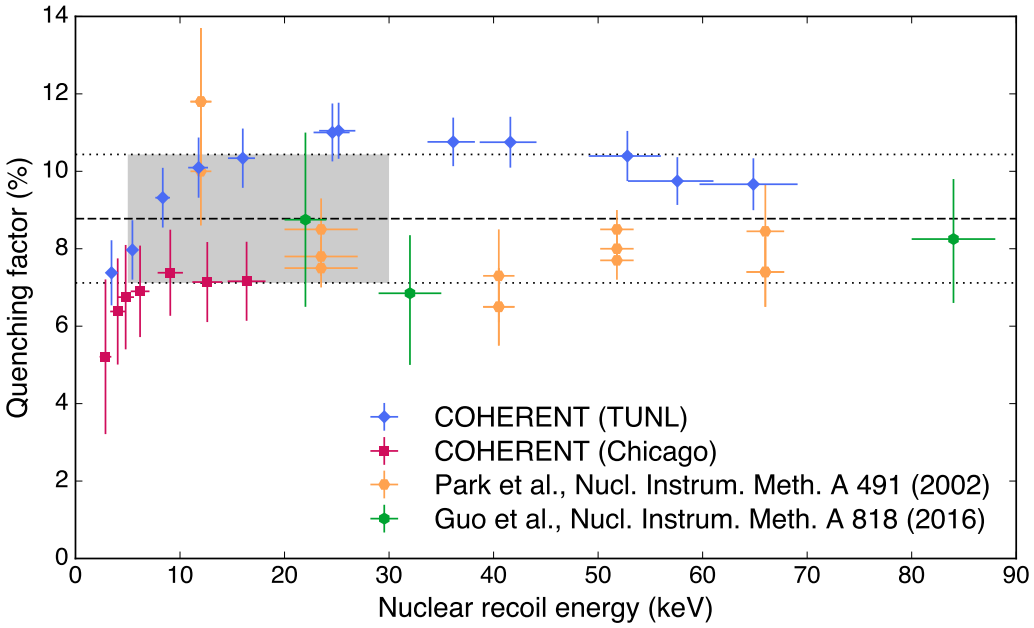

Figure 4.23: Nuclear-recoil quenching factor measurements for CsI[Na] including new data from COHERENT groups at Chicago (crimson squares) and TUNL (blue diamonds) along with literature values from Guo *et al.* [126] and Park *et al.* [195]. The global best-fit QF value that is adopted for the present CsI[Na] CEνNS analysis is shown as a dashed, horizontal black line. The uncertainty used for the global fit is depicted as a shadowed region over the range of data which informs this fit, nuclear recoil energies between 5 – 30 keV. A version of this figure appears in Reference [17].

CHAPTER 5: CsI[Na] CEνNS measurement at the SNS with COHERENT

Section 5.1: University of Chicago CsI[Na] detector

The first CEνNS detection system deployed to the SNS as a part of the COHERENT effort is a CsI[Na]-based system assembled by a group led by Prof. Juan Collar of the University of Chicago and discussed in Refs. [17, 77]. This detector is based on a 14.6-kg CsI[Na] crystal grown by Amcrys using low-radioactivity salts. A low-background OFHC copper can, lined with a PTFE reflector, contains the crystal and a synthetic silica window separates the crystal from the Hamamatsu R877-100 photomultiplier tube.

The R877-100 PMT has a "super-bialkali" photocathode which has increased quantum efficiency (QE) relative to more common bialkali photocathodes (peak efficiencies are roughly 30% and 20%, respectively [77]). As a large fraction of CEνNS events is contained at low nuclear-recoil energies, and correspondingly low scintillation-photon numbers, an increase in QE provides a meaningful increase in the expected number of detectable CEνNS events.

5.1.1: Evaluation and testing at University of Chicago

Prior to installation of the CsI[Na] detector at the SNS, the system was evaluated thoroughly at the University of Chicago by Prof. J. Collar and Dr. B. Scholz. These tests included calibrations using small-angle Compton scattering of low-energy γ-rays from ^{133}Ba which enabled development of waveform-analysis cuts tailored for the CEνNS ROI in photoelectron space. Data was also collected using a ^{241}Am source positioned at several locations along the length of the cylindrical crystal: these data sets were used to produce both a light yield calibration for the crystal, 13.35 PE/keVee, and an assessment of the light-collection uniformity which was found to deviate at most by $< 2\%$. Detailed discussions of all tests are included in Refs. [17, 224, 225].

Section 5.2: Backgrounds at the SNS

Any effort to observe the CEνNS process will share many common background concerns with rare-event searches such as attempts to observe neutrinoless double-beta decay and dark-matter interactions with nuclei. Owing to the significant interest and investment (both capital and intellectual) in rare-event searches, an

understanding of the nature of these backgrounds and techniques for mitigating their effects is well established in the literature [110, 132]. Additionally, due to the proximity to the SNS neutron-production target which is an intense source of both high-energy neutrons and moderate-energy neutrinos, there are some unique background concerns that must be considered.

5.2.1: Steady-state ambient and cosmogenic backgrounds

Though this category of backgrounds represents a wide variety of sources, the common feature is that they show no correlation with the SNS facility operation or the beam. This category encompasses radioactivity from decay or decay-related sources external to the detector; radioactivity from the same sources *within* the detector; cosmic ray interactions within the detector; daughter particles resulting from cosmic ray interactions within the closest layers of shielding; and daughter particles from cosmic ray interactions with external materials or outer shielding layers.

These backgrounds are generally common to those of concern to other sensitive experiments and many of the same mitigation techniques can be employed [110, 132]. A high-efficiency muon-veto system, consisting of 5-cm thick plastic scintillator panels, located inside of the outer-most shielding layer but surrounding the high-Z shielding, allows for the identification and rejection of events which are possibly contaminated by μ-induced neutron production in the experiment shielding [17]. The passive components of the shielding structure, which is discussed in greater detail in Sec. 5.3, work to greatly reduce the ambient backgrounds through a combination of hydrogenous materials (efficient at moderating energetic neutrons and attenuating the total transmitted neutron flux) and lead, a high-Z material which is effective in moderating the flux of γ-rays from external sources.

5.2.2: Prompt neutrons from SNS

The SNS is described as ".. a one-of-a-kind research facility that provides the most intense pulsed neutron beams in the world.." [91], and while the ancillary neutrino output makes a CEνNS measurement realizable, the neutrons produced by the SNS could undermine the viability of any such attempt. Neutron backgrounds are of particular concern for a CEνNS measurement as energy depositions from both CEνNS and elastic neutron scattering are in the form of recoiling nuclei. This concern is shared with WIMP dark-matter searches [116], noting again the common observable in CEνNS interactions and WIMP-nucleus scattering (§1.6, §2.2).

Prompt neutrons incident on the CsI[Na] detector will have traveled out of the liquid-mercury spallation production target, through any close-proximity shielding around the target, and through $\sim$19 meters of a

combination of structural materials (e.g., concrete, rebar) and "backfill", itself a combination of dirt, gravel, or rocks. At the site of the CEνNS-search, flux of low-energy neutrons from the spallation target will be strongly attenuated by the considerable amount of intervening material. However, energetic neutrons which do approach neutrino alley may interact with material near the experiment, either in the shielding, walls, or floor, for instance, and ultimately reach the CsI[Na] installation as lower energy neutrons. Prompt-neutron transport is difficult to model *a priori* with high accuracy.As an approximate model of the spectrum of downscattered neutrons, a power-law energy spectrum is adopted, as suggested by the spectral data from various other neutron-background measurements conducted by the COHERENT Collaboration [17], and determination of parameters for this model is discussed in Sec. 5.5.

Prompt, SNS-beam neutrons produce a background which is fundamentally very similar to ambient or cosmogenic neutron backgrounds. Though the involved energies may differ, the hydrogenous shielding elements (§5.3) should be effective to reduce the flux of prompt neutrons witnessed by the CsI[Na] detector. Since the flux is not reduced by shielding to entirely negligible levels (§5.5), a prompt-neutron feature is included in spectral models used to analyze the collected CEνNS data (§6.2.2).

5.2.3: SNS-operation-related backgrounds

Located within the area of the SNS in which the COHERENT suite of experiments are located (§2.7) is a large pipe that is a part of the SNS building and target exhaust system. This "hot off-gas" (HOG) pipe contains the flow of numerous radioisotopes produced in either the target or other systems; target gases are sent through numerous purification systems to remove mercury, noble gases, and tritium prior to exhaust through the HOG, while other systems are not subject to the same kind of prefiltering [158, 210]. With the target off-gas contaminants mitigated, much of the activity in the HOG is likely from the water cooling system [28, 158], which suggests the HOG contents feature common oxygen activation products such as ^{11}C, ^{13}N, and ^{15}O, all of which decay via β^+ emission, yielding, ultimately, 511-keV γ rays. Though some 511-keV γ flux is expected from ambient sources, the intensity of this additional source has modest variability in time and is confined to a single, extended geometry. Dosimeters situated along neutrino alley, as a part of the radiation safety mission at SNS, observe meaningful γ-radiation dose increases during periods of SNS beam operation, and this is attributed to the HOG.

High-Z shielding meant to address generic γ-ray background sources is a standard component of sensitive-experiment shielding configurations [110, 116, 132]. As the 511-keV γ-rays from the HOG are of modest energy, the contribution from this source to collected spectra in the CEνNS search is easily mitigated and does not require special consideration or unique shielding requirements. The muon-veto systems for experiments in

neutrino alley could be sensitive to HOG γ-rays, however, and some consideration must be given when setting thresholds so that veto-system efficiency isn't meaningfully compromised while maintaining a manageable trigger rate during SNS operation.

5.2.4: Neutrino-induced neutrons

As the most intense pulsed neutrino source in the world [99], the SNS environment is subject to a somewhat unique potential source of background events: neutrino-induced neutron production on materials near detection elements. Neutrino-induced neutrons, or NINs, were considered as a possible background during the LSND experiment [155] and have been offered as a possible explanation for the annual-modulation observed by the DAMA/LIBRA dark matter experiment [84]. Though the NIN-based explanation of the DAMA/LIBRA phenomenon was quickly and soundly refuted [33, 46], the original suggestion nevertheless highlights the need to consider this source of background for very-sensitive experiments, especially in the presence of a neutrino flux more significant than that from the sun.

Section 5.3: Shielding structure and deployment of the experiment to SNS

Shielding of backgrounds was accomplished with a multi-layer, multi-material structure designed by the Collar group [17]. Around the detector itself, the layers of material, proceeding from the inner-most to outer-most, are:

- 7.5 cm of HDPE intended to mitigate neutron backgrounds, particularly contributions from NINs emitted from the surrounding high-Z shielding.
- 5 cm of ancient, low-activity ($\sim$10 Bq ^{210}Pb / kg) lead [17, 77].
- 10 cm of contemporary lead.
- 5-cm thick plastic scintillator panels composing a muon-veto system.
- Water tanks, adding $\sim$9 cm of hydrogenous outer shielding to reduce external neutron backgrounds.

The structure thusly described sits on top of 15 cm of HDPE and is topped with a 5-cm thick plastic-scintillator panel and, suspended further above this panel, $\sim$9" of water shielding. More detailed discussion can be found in Ref. [224]. Figure 5.1 shows side-on and axial views of an MCNPX geometry used in simulations of the CsI[Na] detector system (§5.6). A photograph taken during the installation process at the SNS is shown in Figure 5.2.

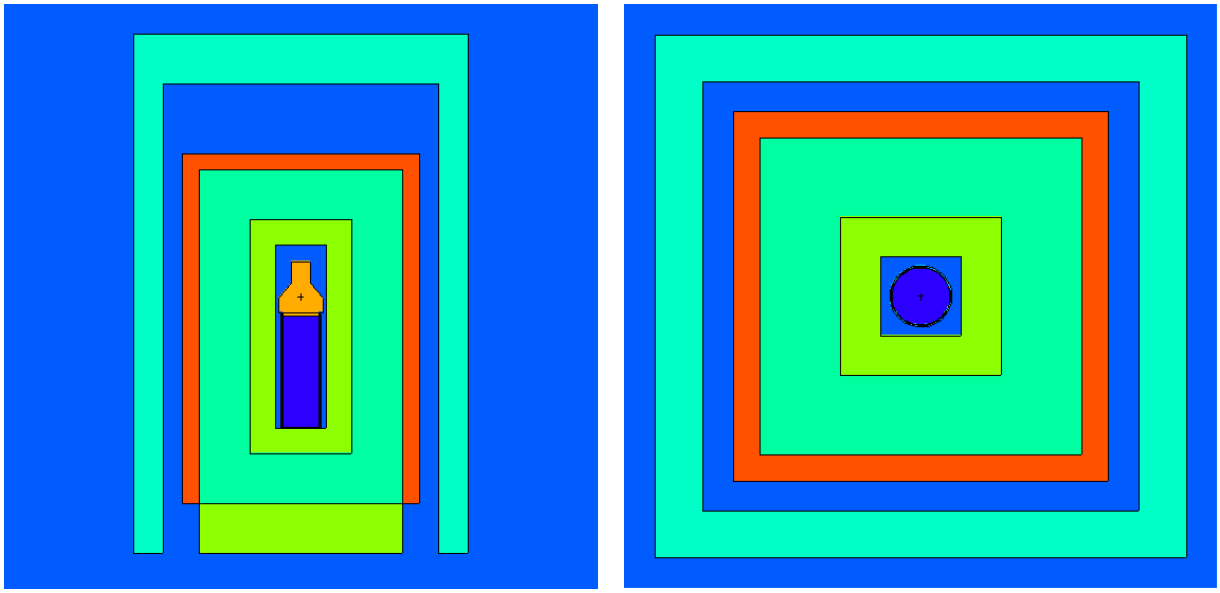

(a) Side view (b) Axial view

Figure 5.1: MCNPX-PoliMi geometry of CsI[Na] detector system installed at the SNS as a part of the COHERENT Experiment. Cross-sectional views are shown from the side (a) and along the vertical axis (b). This geometry can be compared with the partially assembled shielding structure visible in Figure 5.2. Simulation geometry developed by J.I. Collar.

Section 5.4: Data acquisition system and trigger from SNS timing system

An amplified version of the CsI[Na] PMT signal was digitized at 500 MS/s using an 8-bit, 2-channel PCI digitizer made by National Instruments [187, 225]. During testing, it was found that large-integral signals from the PMT, associated with energy depositions in the CsI[Na] crystal which are above the ROI for CEνNS signals, could lead to unstable behavior of the digitizer system so a linear-gate-based circuit was devised to suppress high-energy CsI[Na] events [225]. The second channel of the digitizer recorded an analog voltage signal which took on discrete voltages that indicated the multiplicity level of events in the muon veto panels, allowing for the identification of events which may have μ-related signals present.

A diagram of the DAQ electronics can be seen in Fig. 5.3. This system was designed and implemented entirely by the University of Chicago research group [224, 225].

The digitizer was triggered externally by a signal from the SNS timing system, in which it is referred to as "Event 39". Event 39 is a hardware-generated, time-critical signal in the SNS operation and corresponds to the "extract" signal for the "kicker" magnet system in the proton accumulator ring [54]. The extract signal serves to synchronize the kicker magnets (which act to send the proton beam to the spallation target) and the neutron choppers, which are large, inertial components with very-slow response times; as a consequence of the involvement of the neutron choppers, a primary consideration of the entire SNS timing system is the

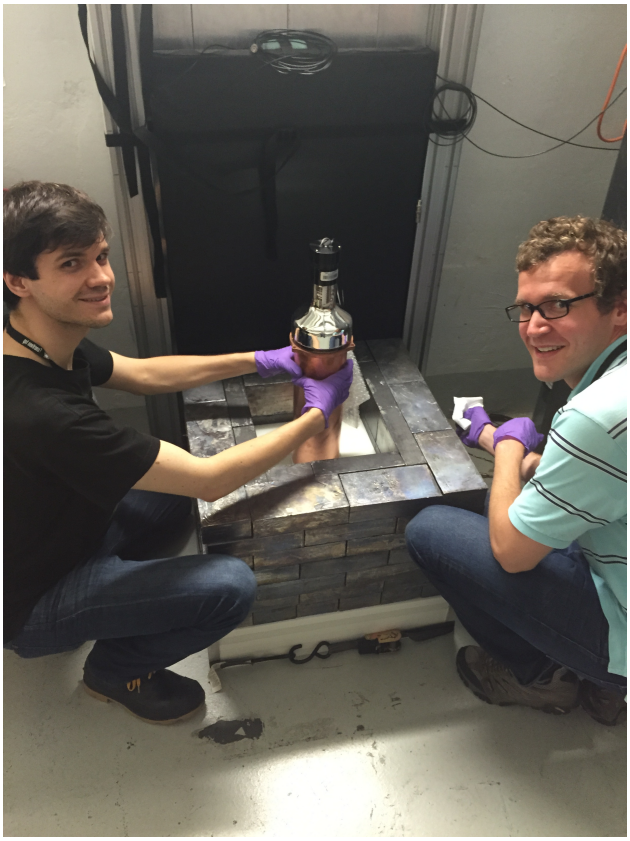

Figure 5.2: Photograph taken during the installation of the 14.6-kg CsI[Na] detector at the SNS, showing Bjorn Scholz (L) and the author (R). Lead bricks are visible, stacked in a staggered manner to mitigate "streaming" of γ-rays into the assembly through interstitial space; the inner layer of lead, with bricks oriented so that less area is visible in the photo, is constituted from low-activity lead, reducing the exposure to the detector of γ-rays from ^{210}Pb decay. The detector is visible in the center of the partially-assembled lead shielding structure, supported by Dr. Scholz. The detector is resting upon HDPE which serves to shield against neutron backgrounds. Photo credit: Juan Collar.

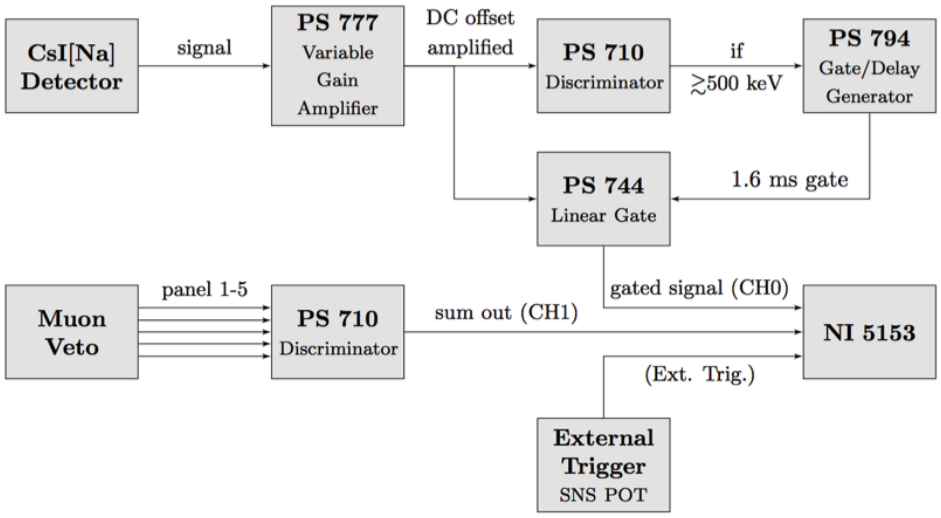

Figure 5.3: Diagram of the data acquisition (DAQ) system in place at the SNS for the CsI[Na] CEνNS experiment. A 500-MS/s digitizer records two waveforms: the signal from the PMT observing the CsI[Na] detector, after it is sent through an amplifier and a linear gate, the latter of which protects the digitizer system from high-charge signals; and a discrete-level analog signal which represents the number of active panels in the muon-veto system. The common trigger for both channels of the digitizer is an external signal, labeled "SNS POT" in the diagram, from the SNS timing system. This trigger signal is "the most time-critical event" in the operation of the SNS [54] and should provide a trigger that is very-stable in time with respect to neutrino production over long periods of operation (see Section 5.4 of the text). The DAQ was designed and put in place by the research group led by Prof. Juan Collar at the University of Chicago; for more discussion, see Reference [224]. Figure from Reference [225].

stability of the extract signal [54].

As Event 39 is related to the beam delivery, the relative time between its arrival and the incidence of the proton beam on the mercury target (this relative time will be referred to as POT offset) should be very stable over the course of an SNS "run": a period of stable beam production during which beam energy is not changed, typically several weeks or months long. Different SNS runs may utilize different proton beam energies and small differences in the POT offsets between runs may be observed. The significance and determination of the POT offset is discussed in §5.5.3.

Section 5.5: Measurement of neutron backgrounds

An understanding of neutron backgrounds is of crucial importance for a successful observation of CEνNS and separate experimental, simulation, and analysis efforts were dedicated to this subject. To enable a robust, multi-dimensional analysis of the CEνNS search data, information on the spectral and timing distributions of neutron backgrounds was sought. These efforts were focused on the unique problems presented by the SNS environment in the form of prompt, SNS-beam neutrons (§5.2.2) and neutrino-induced neutrons (§5.2.4).

In the case of prompt neutrons, the spectrum incident upon the detector assembly is assumed to be modeled by a power-law distribution,

$$f(E_n) \sim E_n^{-\alpha}, \tag{5.1}$$

with the parameter α governing how "hard" the spectrum is: lower values of α yield harder spectra with substantial contributions from higher-energy neutrons. With these background measurements, we seek to determine a value for α that effectively describes the energy spectrum of the SNS-produced neutrons which travel from the mercury target, through the intervening material, and impinge on the shielding structure. This value of α will then be used as an input to simulations of the CEνNS search geometry, yielding a model for the recoil-energy distribution we would expect to be contributed by this background component.

Complementing the energy spectrum determination, limits can be placed on the rate of prompt-neutron events that should be seen in the CEνNS data and the arrival time of these neutrons with respect to the SNS beam timing pulse can be established. With these measurements, not only can the prompt-neutron background model can be well-constrained, but the arrival times for beam-related signal and background components can be established for any CEνNS search.

5.5.1: Advance deployment of liquid scintillator cells in shielding structure

Prior to the installation of the CsI[Na] detector, an effort was made to measure neutron backgrounds of the intended CEνNS search. These measurements were carried out by installing two 1.5-L liquid scintillator cells at the eventual location of the CsI[Na] detector in the SNS basement. The liquid scintillator cells used Eljen EJ-301 scintillator and fast PMTs to maximize pulse-shape discrimination capabilities [17].

Though the location of this measurement was the same as the CsI[Na] detector, the shielding structure was *not* entirely shared between the two experiments. To perform the neutron background measurement, a shielding structure planned in Ref. [77] was constructed[1], sitting atop a ∼6"-tall pedestal of HDPE serving to attenuate neutron flux, especially from the concrete floor. The outer layer of hydrogenous material, consisting of HDPE planks with a total thickness of ∼3 cm, was also intended to reduce the flux of external neutrons. Directly interior of the planks were 5-cm-thick muon-veto panels which allow the rejection of events during which a cosmic ray was incident upon the system; composed of plastic, these panels also serve as supplemental neutron shielding. The inner-most layer of shielding for the neutron background measurement was composed of lead constructed in a hexagonal pattern such that the minimal thickness surrounding the inner cavity was ∼15 cm. This minimum thickness was accomplished both radially and axially: the lead

[1]The description in Ref. [77] pertains to the deployment of the CsI[Na] detector and the structure built and described here, intended to house the liquid scintillator cells, deviates slightly from the original prescription. Specifically, the structure used in the measurement of neutron backgrounds does not include an inner-most, 1"-thick layer of low-background lead.

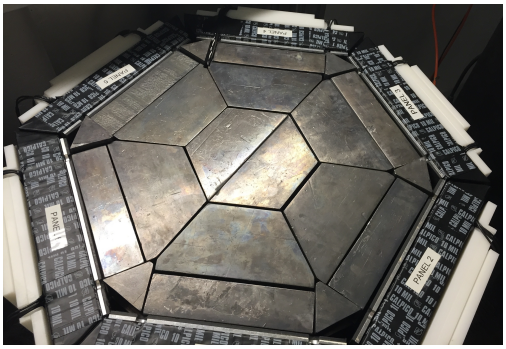

(a) Top of shielding assembly for *in situ* neutron background measurement for CsI[Na] detector.

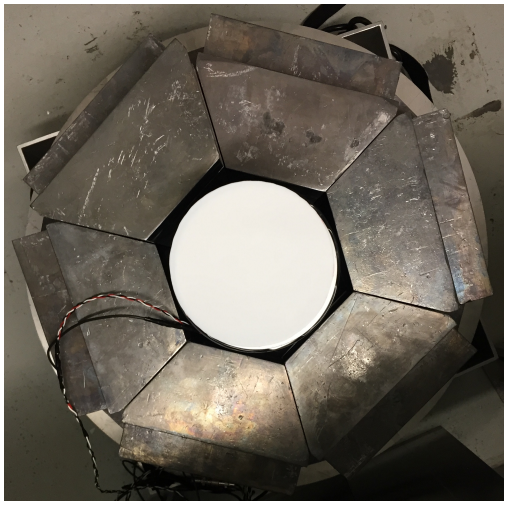

(b) Top-down view of partially-assembled structure for *in situ* neutron background measurement. The bottom liquid scintillator cell is visible in the center.

Figure 5.4: Images of the structure housing the liquid-scintillator cells used for *in situ* measurements of neutron backgrounds. The shielding structure differs slightly from that utilized during the CsI[Na] CEνNS search and is described in Reference [77]. Figure 5.5 depicts the MCNPX geometry modeling this assembly and affords an alternative perspective.

structure had a pedestal and "roof" which were 15-cm thick. A final 10 cm of hydrogenous material sat on top of the structure, with a top-most layer of HDPE supported on a circular, 5-cm muon-veto panel.

Inside of the structure, the liquid scintillators were arranged so that they stood length-wise and the cells met near the center of the cavity. This entire assembly was surrounded with an additional layer of external water shielding which would also be used with the CsI[Na] installation (§5.3). Some photographs of the partially assembled shielding structure and detector configuration can be seen in Fig. 5.4.

5.5.2: Simulation, analysis, and determination of effective neutron-background model

To establish an effective α parameter and prompt-neutron flux, simulations were carried out using the geometry shown from two perspectives in Fig. 5.5 with numerous initial distributions of the neutron energy. This geometry is representative of the configuration described in Sec. 5.5.1 and was coded by J.I. Collar. Neutron energies were described by power-law distributions with different values of α, ranging from 0 to 2 in steps of 0.025. Figure 5.6 shows several distributions with different α values.

The output of these simulations represents a record of energy depositions in the liquid scintillator cells located at the center of the shielding structure[2]. These histories were converted into electron-equivalent energies in a process similar to that described in Sec. 4.4.2, accounting appropriately for the distinctions

[2]Note that the analysis of the simulation output was carried out by P.S. Barbeau but has been summarized here for context.

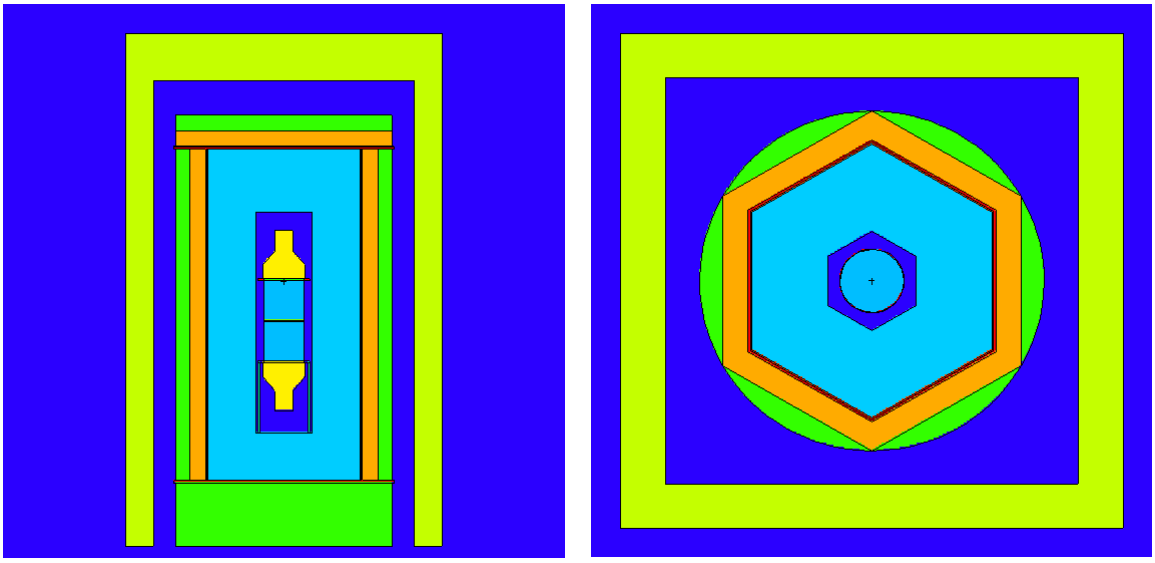

(a) Side view (b) Axial view

Figure 5.5: Cross-sectional views of the geometry for MCNPX-PoliMi simulation of *in situ* neutron background measurements. Panel (a) depicts a side view and (b) shows an axial view. The impression of a rectangular geometry in the side view is an artifact of the cross-sectional perspective. Photographs in Figure 5.4 depict the actual experiment modeled here. Simulation geometry developed by J.I. Collar.

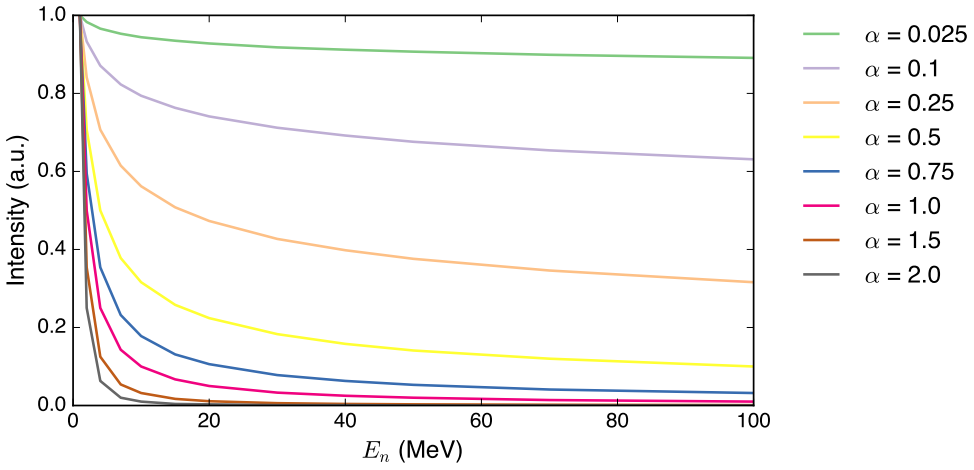

Figure 5.6: Power-law neutron-energy distributions used in the simulation of the prompt, SNS-beam neutron background. Simulations were carried out with α-parameter values ranging between 0 and 2 in 0.025 increments. The simulated interaction spectra were compared against experimental data to determine a value of α which effectively models the spectrum of SNS neutrons incident on a detector at this location. Distribution values from P.S. Barbeau.

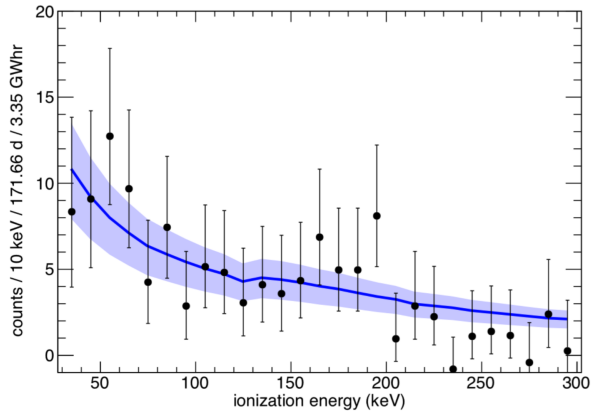

(a) Observed ionization energy in liquid scintillator cells measuring neutron background. The blue line shows the spectrum for the best-fit values of α and flux for the prompt neutron background; the shaded region is the 1σ uncertainty band.

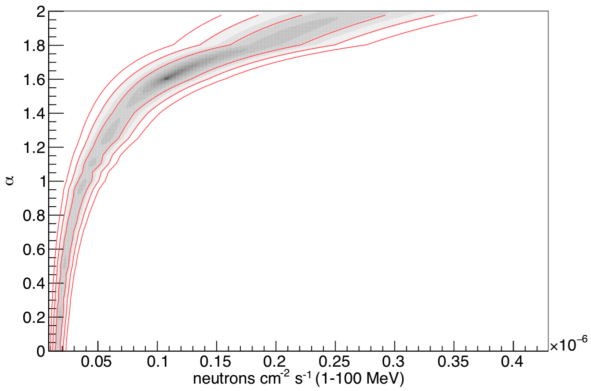

(b) Likelihood map for fits to collected prompt-neutron-background data with different values of spectral-parameter α and prompt-background flux.

Figure 5.7: Results of fitting the *in situ* neutron measurements with a power-law neutron spectrum: (a) shows the data from the EJ-301 liquid scintillator cells with the best fit neutron spectrum overlaid, representing both the flux normalization and the α spectral parameter; (b) depicts the likelihood values at different values of α and fluxes. Analysis by P.S. Barbeau. From The COHERENT Collaboration [17].

between depositions from γ-rays/electrons and neutrons, and taking into consideration the nuclear-recoil-energy dependent light yield of the EJ-301 scintillator [17]. A normalized distribution of observable energy depositions was thus constructed for the neutron spectra associated with each value of the spectral parameter α.

Experimental data representing 171.66 days of collection and an SNS-beam exposure of 3.35 GW-hr was fit with the simulation-produced distributions. The overall normalization of the distribution was allowed to float and the fitted value determined a neutron flux (in units of neutrons / cm^2 / s) in the neutron-energy interval between 1 and 100 MeV. For each value of α, ranging between 0 and 2 in steps of 0.025, a likelihood scan was carried out in the region surrounding the best-fit flux normalization and the likelihood values were recorded for each value of α at fine steps of total flux. Figure 5.7b shows the map of likelihood values. The likelihood is maximized with a spectral-parameter value of $\alpha = 1.6$ and a flux of 1.09×10^{-7} n/ $\left(\text{cm}^2 \cdot \text{s}\right)$ [17]. Collected data from the liquid scintillator cells can be seen in Fig. 5.7a, with the best-fit distribution and 1σ uncertainty band overlaid [17].

Figure 5.8 presents an alternative view of the data collected from the advance deployment of liquid scintillators which assists in the visualization the different components participating; the blue dotted line in this figure depicts the fitted contribution from NINs, which underlies the most prominent feature of the spectrum, itself associated with prompt neutrons. Based on the analysis of this data, the NIN contribution is small ($\sim 1/2$) compared to that of the prompt neutrons [17], the NIN component is not included in

subsequent stages of the analysis.

5.5.3: Confirmation of trigger timing offset

The final, important piece of information which can be extracted from the *in situ* neutron measurement is associated with the timing of signals. As described in Sec. 5.4, a facility-wide timing signal is distributed at the SNS and this is used to trigger the data acquisition system. What is unknown *a priori* is the magnitude of any offset in time which may exist between the arrival of the trigger signal at the DAQ and the arrival of SNS protons on the mercury target. The CEνNS search and the *in situ* neutron measurement share an identical trigger configuration, presenting an opportunity to calibrate a timing offset based on the *in situ* measurement.

From simple kinematic considerations, it is clear that, for particles produced by the same POT pulse, neutrinos should arrive at the CsI[Na] detector prior to beam-related neutrons. Relative timing with respect to the POT signal that triggers the digitizer is not known *a priori* due to lack of knowledge of the cable lengths at various stages, transducer response, and any propagation delay that may be introduced by any involved electronics.

GEANT4 Monte Carlo simulations of neutron production and transport at the SNS, carried out by COHERENT Collaborators[3], had provided preliminary values for the timing offset between the POT signal and neutron arrival at the detector [17], see the inset in Fig. 5.8. It was found that the *in situ* data agreed well with the simulated timing distributions, so the simulated offset was adopted for subsequent stages of the analysis. The data from the measurement with liquid scintillators can be seen in Fig. 5.8 along with the fitted model, representing a uniform background of accidental coincidences, NINs, and prompt neutrons. The adopted offset between the POT DAQ trigger was 1.19 μs, and based on simulations of both the SNS beam neutrons and the SNS neutrinos [213], the relative timing offset between the arrival of the prompt neutrino population and the prompt neutron population was determined[4] to be 0.084 μs.

Section 5.6: Simulations of CsI[Na] detector system

To inform the numerous analyses which must be carried out in support of the CEνNS search, several different Monte Carlo simulations were carried out, including the geometries of both the liquid-scintillator measurements (§5.5.1) and the CsI[Na] experiment itself (§5.3). These simulations focus exclusively on

[3]These GEANT4-based simulations were carried out by G. Perumpilly of the University of Chicago; they are mentioned briefly in Ref. [17].

[4]These timing analyses are attributable to P.S. Barbeau.

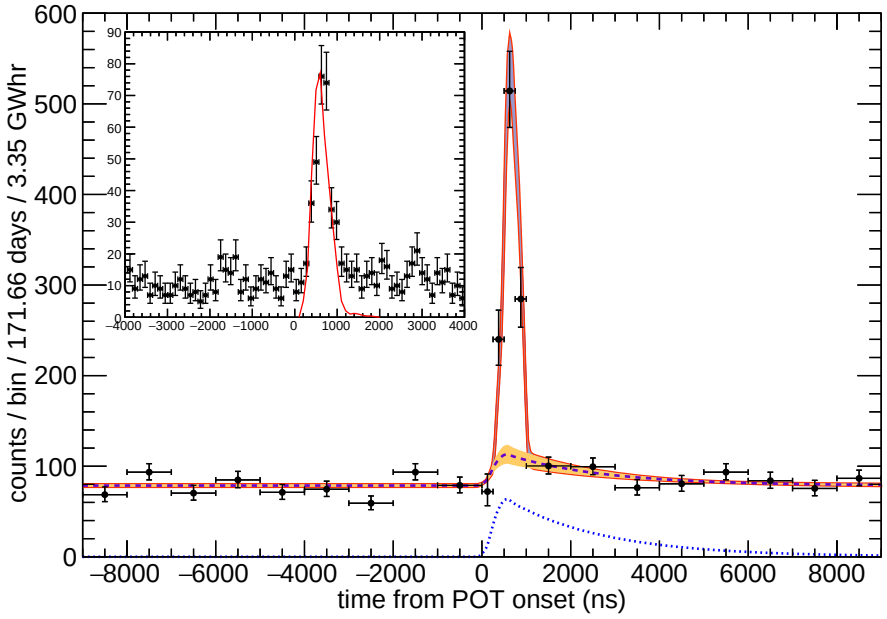

Figure 5.8: Timing data collected during the *in situ* neutron background measurement at the SNS. Overlaid is a model consisting of a uniform background of accidental coincidences, a prominent peak associated with SNS neutrons, and a smaller feature associated with neutrino-induced neutrons (shown as a dotted blue line). Analysis by P.S. Barbeau; from The COHERENT Collaboration [17].

neutron backgrounds and are based on MCNPX-PoliMi; γ-ray backgrounds are accounted for entirely by background data, not coincident with the SNS beam, and no simulations are necessary.

Once appropriate model parameters describing the prompt neutron background component were determined (§5.5), they were used to simulate this background component in the CEνNS-search geometry. Using the geometry shown in Fig. 5.1, which closely approximates the configuration realized at the SNS, the prompt neutron spectrum was simulated. Neutrons were produced uniformly on a circular surface with a radius of 1 m, located $\sim$50 cm from one of the sides of the shielding structure, and oriented vertically. Drawing from a power law distribution with $\alpha = 1.6$, based on the results of Sec. 5.5, 1×10^7 neutrons were generated with their initial direction of travel perpendicular to the face of the source disk and in the direction of the shielding structure. These neutrons effectively illuminated the side of the shielding structure facing the source. The simulation was configured to produce an MCNPX-PoliMi tracking output ("DUMN1") which registered information about interactions in the CsI[Na] volume for any event which possessed interactions therein.

Section 5.7: CEνNS recoil rate calculations

Calculation of the expected recoil rate for CEνNS signals starts with an expression for the differential cross section of the process, incorporates the detector response, and yields an expected number of signal events per unit energy per unit beam exposure which is subsequently used in analysis of the collected data. By producing a recoil spectrum, rather than simply a number of counts above threshold, energy and timing information of the signal can be incorporated into a more-powerful analysis of the data.

For the analysis presented in Chapter 6, the recoil distributions were produced by P.S. Barbeau, but the discussion here follows an analogous procedure.

The differential CEνNS cross section of Eq. (1.1) combined with the neutrino energy distributions for the 3 flavors produced by the SNS (§2.5.2) yield the distribution of events as a function of nuclear recoil energy. For the present analysis, many parameters which go into Eq.(1.1) are fixed for both computational simplicity and for negligible contributions to shape or rate differences.

5.7.1: Form factors

Evaluation of the differential cross section for CEνNS, Eq. (1.1), requires the nuclear form factor $F\left(Q^2\right)$. Inclusion of the form factor, which is a Fourier transform from real space to momentum space of the nuclear density distribution and thus is related to the physical *size* of the nucleus (see §1.4), accounts for loss of coherency at higher values of momentum transfer [170]. Nuclear density is often modeled by the Woods-Saxon [152] or Fermi [170] distributions, neither of which yields an analytic form factor. There are numerous examples in the literature[5] of tractable nuclear form factors which are based on alternative, but still-realistic models of nuclear distributions; we restrict our discussion to two such models.

By treating the nucleus as a solid sphere "folded" with an exponentially decaying radial distribution representing the diffuse nuclear surface, as proposed by Helm [130], one arrives at the so-called Helm form factor[6] given by [101]

$$F\left(Q^2\right) = \frac{3j_1\left(QR_0\right)}{QR_0} \exp\left[\frac{-\left(Qs\right)^2}{2}\right], \tag{5.2}$$

where $j_1(\ldots)$ is a spherical Bessel function of the first kind. In Eq.(5.2), R_0 is described by $R_0^2 = R_{\rm nuc}^2 - 5s^2$, where $R_{\rm nuc}$ is the radius of the nucleus and s is its skin thickness.

[5] CEνNS-specific literature contains some discussion on the subject of form factors (see, e.g., Refs. [136, 221]), but the community focused on direct detection of dark matter has also produced a large body of work on the subject: examples include Ref. [101], devoted to the subject of form factors, and Ref. [170] which includes a review of form factor treatments among many other topics.

[6] When evaluating expressions for the form factor (i.e., Eq. (5.2) or (5.4)), be aware that natural units are used. If one is working in units of fm, keV, s, then to use Eqs. (5.2) and (5.4) simply express Q in keV (dropping the $1/c$) and R in fm; then, wherever the expression QR appears, include a factor of $(\hbar c)^{-1} \approx 197332$ keV $\cdot$ fm.

Nuclear radii are traditionally approximated in terms of the mass number A as [162]

$$R_{\text{nuc}} \approx 1.2A^{1/3} \text{ fm}, \tag{5.3}$$

and Engel takes $s \sim 1$ fm [101]. Calculations shown here for the Helm form factor adopt the same values as Ref. [101].

Klein and Nystrand [152] propose approximation of the nuclear potential as a hard sphere of radius R_A convolved with a Yukawa potential of range $a = 0.7$ fm. This approach yields a form factor given by [152]

$$F\left(Q^2\right) = \frac{4\pi\rho_0}{AQ^3}\left[\sin\left(QR_A\right) - QR_A\cos\left(QR_A\right)\right]\frac{1}{1+a^2Q^2}, \tag{5.4}$$

where R_A is the traditional approximation of the nuclear radius $R_A \equiv R_{\text{nuc}}$ of Eq.(5.3) and ρ_0 is a normalization parameter such that $F(0) = 1$. To maintain consistency with Klein and Nystrand, the numerical prefactor of 1.2 in Eq.(5.3) is unchanged, in contrast with the treatment used here for the Helm form factor.

The momentum transfer Q which appears in expressions for the form factor $F(Q)$ is expressed most generically as [136]

$$Q^2 = \frac{2E_\nu^2 E_{\text{rec}} M}{E_\nu^2 - E_\nu E_{\text{rec}}}, \tag{5.5}$$

in terms of the energy of the incident neutrino E_ν, the energy of the recoiling nucleus E_{rec}, and the mass of the nucleus M. In the case of CEνNS, we can quite reasonably make the approximation that $E_{\text{rec}} \ll E_\nu$, so we rewrite (5.5) in terms of the quantity $E_{\text{ratio}} = E_{\text{rec}}/E_\nu$ and expand about $E_{\text{ratio}} = 0$, finding

$$\begin{aligned} Q^2 &= \frac{2E_\nu^3 E_{\text{ratio}} M}{E_\nu^2 - E_\nu^2 E_{\text{ratio}}} \\ &\approx 2ME_\nu E_{\text{ratio}} + 2ME_\nu E_{\text{ratio}}^2 + 2ME_\nu E_{\text{ratio}}^3 + \ldots \\ &\approx 2ME_{\text{rec}} \\ Q &\approx \sqrt{2ME_{\text{rec}}}. \end{aligned} \tag{5.6}$$

Keeping only the first term in the expansion of Q^2 simplifies subsequent computation (removing the dependence of Q^2 on an additional parameter, E_ν) with little effect on the precision of the approximation: characteristic neutrino energies for the SNS are $\mathcal{O}(10$ MeV) while recoil energies are $\mathcal{O}(10$ keV).

Figure 5.9a shows the Helm and Klein-Nystrand form factors for a nucleus of mass $A = 133$ in terms of Q; these same form factors are shown in terms of E_{rec} in Figs. 5.9b and 5.9c. The approximate relationship between E_{rec} and Q, given in Eq.(5.6) and used to calculate the form factors in terms of E_{rec}, is drawn for

clarity in Fig. 5.9b.

For the recoil shapes used in the subsequent analysis of the CsI[Na] CEνNS search[7], the form factor prescription of Klein and Nystrand [152] is adopted. Shape parameters $a = 0.7$ fm and $R_A = 1.2A^{1/3}$ fm are used, consistent with the original description of the approach [152].

Barranco *et al.* [38] consider the choice of different form factor models and the impact on the sensitivity of CEνNS experiments to new physics. The authors of Ref. [38] employ the Helm model, as presented by Engel [101] and replicated here in Eq. (5.2), but they also consider the form factors of Ahlen *et al.* [10], Freese *et al.* [113].

5.7.2: Recoil rate distributions

The recoil distributions used for the subsequent analysis can be seen in Fig. 5.10; these were calculated by P.S. Barbeau, using Standard Model predictions and the full differential cross section of Eq. (1.1). These rates make use of the Klein-Nystrand form factor.

In the case of the prompt neutron distribution, the shape is informed by an MCNPX-PoliMi simulation of the CsI[Na] assembly geometry with externally produced neutrons using an energy distribution whose parameters were determined in Sec. 5.5.2 based on the *in situ* background measurement.

Since these distributions are in terms of the number of observed photoelectrons, information about the photoelectron yield and quenching factor for CsI[Na] must be utilized in converting the nuclear-recoil distribution described by Eq. (1.1). The utilized quenching factor is that measured in the preceding chapter, and the photoelectron yield was measured prior to deployment of the 14.6-kg CsI[Na] CEνNS detector. Variation of the QF results in change of the recoil distribution shape, but these effects were found to be very small compared to the overall effect on the rate, so only the normalization of the distributions shown in Fig. 5.10 are ultimately varied in the fitting process described later.

[7]Klein-Nystrand is used to produce the shapes presented in this section as well as the recoil distributions, produced by P.S. Barbeau, that were ultimately used in the statistical analysis of Chapter 6.

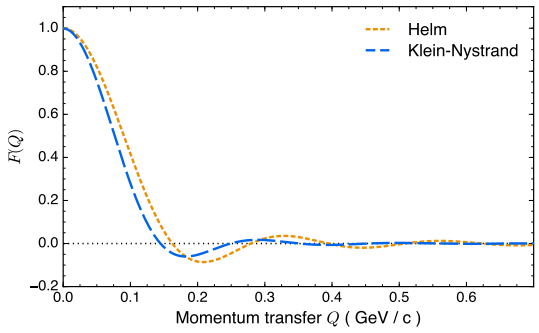

(a) Nuclear form factor from both the Helm and Klein-Nystrand formulations in terms of the momentum transfer Q.

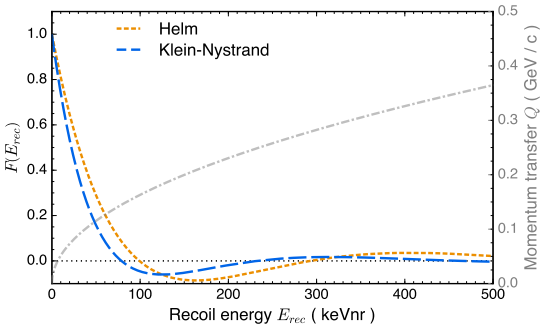

(b) Nuclear form factor in terms of nuclear-recoil energy from both the Helm and Klein-Nystrand formulations. The input to the form factors is an approximation of the momentum transfer Q depending only on the recoil energy; this approximation, given by Equation (5.6), is shown as a dash-dotted, gray line.

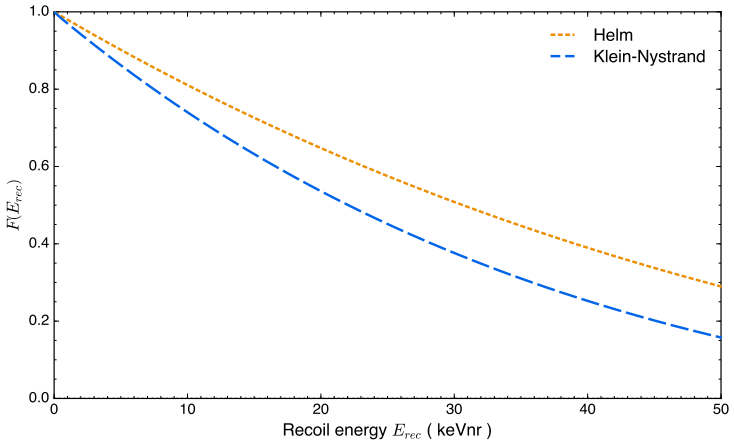

(c) As in (b), this shows the Helm and Klein-Nystrand form factors in terms of E_{rec} using Equation (5.6) to approximate momentum transfer. The plotted region is restricted to the CEνNS ROI to better highlight differences in the form factors relevant to CEνNS detection.

Figure 5.9: Form factors used to describe the nuclear spatial distribution for $A = 133$. The Helm form factor (shown as a dotted, dark-orange line) which arises from adopting the nuclear density description of Reference [130], and the model proposed by Klein and Nystrand [152] (dashed, dodger-blue line), are calculated using Equations (5.2) and (5.4), respectively. The nuclear radius is approximated differently between the two models as drawn: see the text for details. Form factors are shown both in terms of: momentum transfer Q in plot (a); and nuclear-recoil energy E_{rec} in plots (b) and (c). The approximate relationship between Q and E_{rec} is plotted in (b).

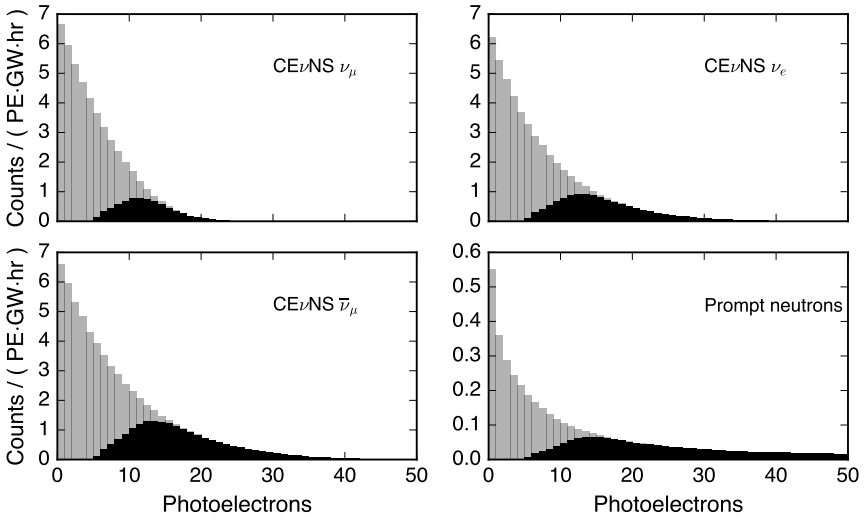

Figure 5.10: Recoil rate distributions for CsI[Na] at the SNS as used in subsequent analysis. The shaded region is the raw recoil rate anticipated based on the simulated neutrino spectrum shown in Figure 2.2b at the appropriate distance from target. For comparison, the "accepted" rate distribution is also included, shown in black; this after applying the UChicago acceptance efficiency to the predicted rates, though an overall scaling factor of 0.6655 is omitted to better visualize when the efficiency plateaus. See Section 6.1.2 for a discussion of acceptance efficiencies. All count rates are normalized to 1 GW $\cdot$ hr of beam exposure.

CHAPTER 6: Statistical analysis of CEνNS search data with a CsI[Na] detector

Section 6.1: Parallel analysis and signal-processing pipelines

The data acquired by the CsI[Na] detector system was processed independently through two separate and well-segregated signal-processing pipelines, which transform the raw waveforms collected by the DAQ into data sets with reconstructed energy and timing information for interactions in the CsI[Na] detector. Groups located at the University of Chicago (UofC) and the combination of the National Research Nuclear University, Moscow Engineering Physics Institute (MEPhI) and the Institute for Theoretical and Experimental Physics (ITEP), part of the Kurchatov Institute, led these analysis efforts, and this work will contribute significantly to the doctoral theses of Bjorn Scholz (UofC) and Alexey Konovalov (MEPhI/ITEP). Brief descriptions of the analysis approaches are included here, while full discussions can be found in [156, 224].

6.1.1: Waveform time windows

The DAQ (§5.4) is configured to record 70-μs waveforms triggered on the SNS timing signal[1]. DAQ settings produce waveforms which begin 55 μs prior to the trigger signal. Two analysis regions are defined within the waveforms [17]: the "coincident" region (C), which begins at $t = 55$ μs; and the "anti-coincident" region (AC) which begins at $t = 40$ μs.

For each analysis region, the preceding 40 μs constitutes a "pretrace" region which is used to reject events that demonstrate considerable evidence for contamination by scintillation afterglow originating from preceding, high-energy depositions in the CsI[Na]. The timing window definitions are shared between the two analysis pipelines but the distinct cuts applied are the result of considerable, independent development by Konovalov [156] and Scholz [224].

6.1.2: Acceptance curves

A key characteristic of the analyses are the acceptance curves, which describe the efficiency for detection of a signal as a function of the signal integral. Acceptance curves for the UofC and MEPhI/ITEP analysis can be seen in Figs 6.1. These efficiencies participate in the determination of the expected spectral characteristics

[1] See Section 5.4 for a discussion of the DAQ trigger signal: "SNS timing signal", used here for clarity and simplicity, refers to Event 39 of the SNS timing system, more specifically the version of this which is distributed to the fiber-optic timing decoder located in "neutrino alley".

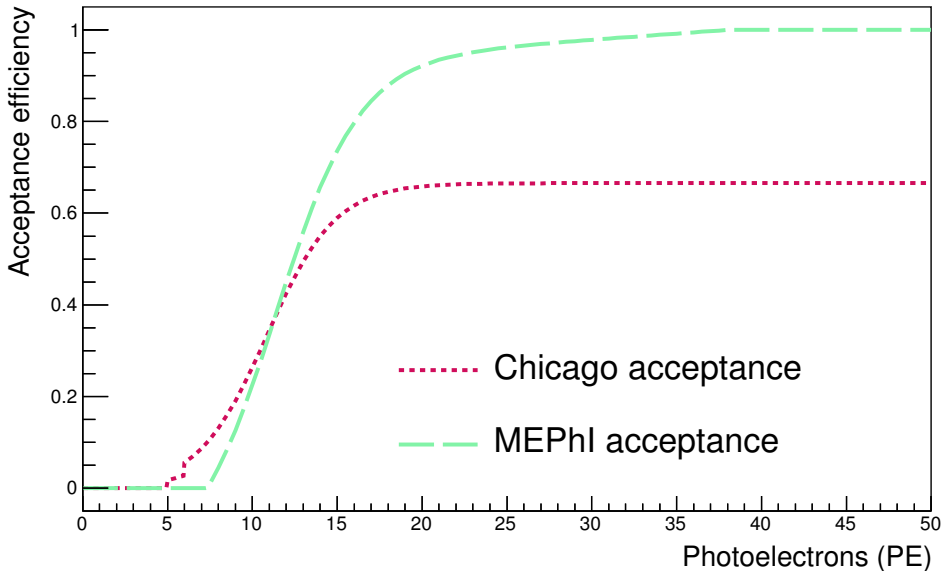

Figure 6.1: Acceptance curves for CsI[Na] events from both analysis pipelines. The Chicago acceptance curve, associated with the analysis by B. Scholz and described more fully in his thesis [224], is shown as a finely-dashed, crimson line; the MEPhI/ITEP acceptance, associated with the analysis by A. Konovalov and described more fully in his thesis [156], is depicted as a coarsely-dashed, light-green line. These curves play critical roles in the development of the expected spectral shapes which define the PDFs used in much of the statistical analysis of the CsI[Na] CEνNS data. The difference in maximal acceptance efficiencies is attributable to different approaches regarding the treatment of live time.

for each analysis pipeline which inform the PDFs used in subsequent analysis; development of these PDFs is discussed in Sec. 6.2.

Comparing the acceptance curves shown Fig. 6.1, one will observe that the efficiency approaches unity for the MEPhI/ITEP analysis but reaches an asymptotic maximal value of 0.666 in the case of the Chicago analysis. This is the result of different approaches for treating live-time effects in the separate analyses; this difference is accounted for in subsequent stages of the statistical analysis.

6.1.3: Common "reduced" analysis output format

Both analyses yield "reduced" data sets which are a collection of "accepted" events in the C and AC regions for the respective beam-on and beam-off periods. These accepted events have effectively been subject to acceptance cuts shown in Fig. 6.1, though the process of reduction involves the application of many well-refined requirements; some discussion can be found in Ref. [17], though full representation of the efforts can be found in the theses of the analyzers [156, 224]. The reduced data sets are those which are used for the statistical analysis of the CEνNS experiment in conjunction with the PDFs informed by the appropriate acceptance curve.

Section 6.2: Spectral features of CEνNS search data

In all but the most simple of approaches (§6.4), analysis of the CEνNS search data requires a model for each of the signal and background components that are present in the collected spectra. Here we will develop *a priori* PDFs to describe the different components; these PDFs will then be used in several different approaches to the analysis presented in subsequent sections.

6.2.1: Steady-state backgrounds

Development of the steady-state background model makes use of data collected in the AC region of the beam-on dataset (§6.1.3). This empirical approach is in part enabled by the long digitized waveforms that are recorded and triggered by the SNS timing signal: as there is no hardware threshold placed on the CsI[Na] signals, for a given event the "signal" region is distinguished from the "background" region entirely by their relationship in time with respect to the SNS beam. Any accidental coincidence should have equal probability of occurring in either the signal or background regions of the waveform, so the steady-state backgrounds (which are uniformly distributed in time) should be represented faithfully by the sample provided in the AC window.

To reduce the effect of statistical fluctuations in the background model, the 2-D AC dataset is not employed directly as the model. Recognizing that the time and energy dimensions of the background data are uncorrelated[2], the AC data is factorized, producing 1-D projections in energy and arrival-time space. With this approach, few-count-per-bin statistical effects are mitigated relative to the 2-D case, where many bins may have zero or $\mathcal{O}(1)$ counts.

The energy projection is taken, unmodified, to represent the energy-space background. To construct the background in arrival-time space, we recall that these steady-state backgrounds should be uniformly distributed in time and consider the expected distribution for the arrival time of these background events, ultimately determining that the analytical form should be that of an exponential decay. An exponential is then fit to the arrival-time projection of the AC data, and the resulting distribution serves as the background model in arrival time. Figure 6.2 depicts the 2-D AC-region data from the Chicago analysis and the projections onto the photoelectron and arrival-time axes. A 2-D steady-state background model is created by taking the product of the photoelectron projection with the arrival-time exponential model. Each photoelectron row of this product is normalized so that the photoelectron projection matches that of the input data. The resulting background model can be seen in the lower right panel of Fig. 6.10.

[2]Prof. J. Detwiler of the University of Washington, a COHERENT collaborator, was the one to make this recognition and suggest the approach that follows.

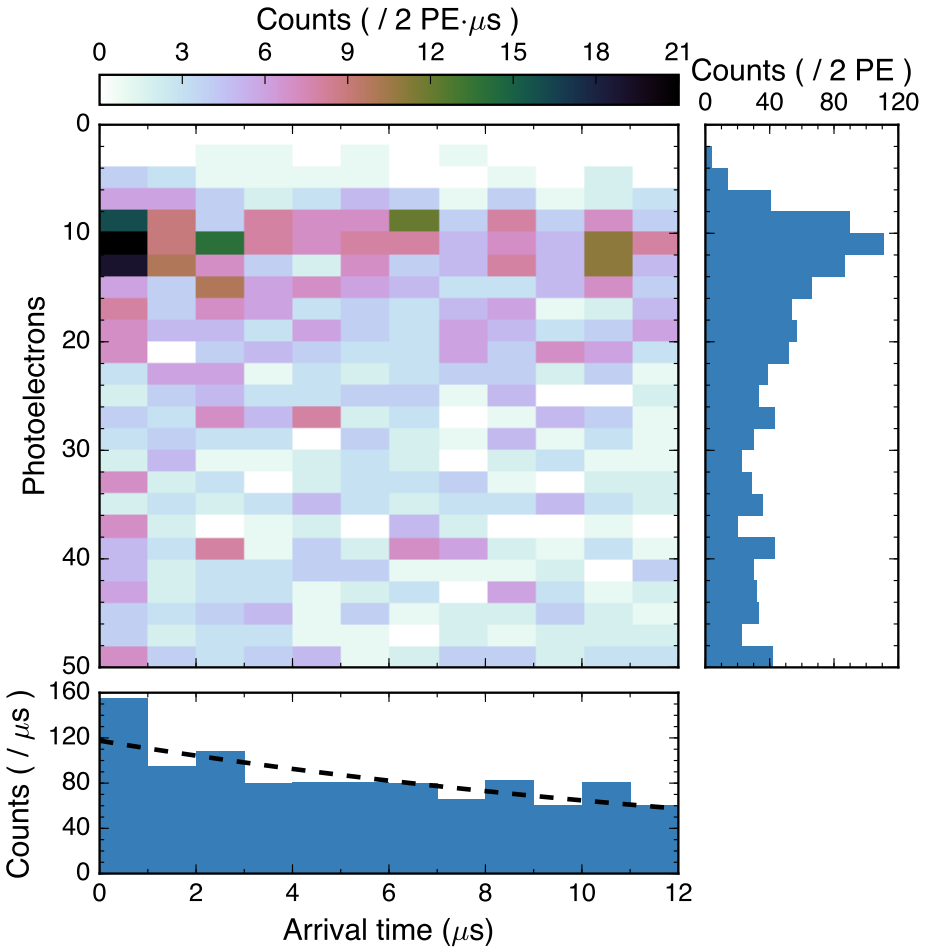

Figure 6.2: Anti-coincidence region data from the CsI[Na] CEνNS search used to inform a model of the steady-state background. In the two-dimensional data, the statistical fluctuations on a bin-by-bin basis complicate direct use of this data as a background model. To construct a two-dimensional model which mitigates statistical effects, a "factorized" construction is employed: the data is projected onto each axis; the photoelectron-space projection is taken, unaltered, and multiplied by a exponential fit to the time-space projection; the resulting product is used as a background model. The projections are shown here alongside the anti-coincident data. The exponential fit to the time projection is drawn on the time-space axis as a dashed, black line. This data and the corresponding model are associated with the Chicago analysis.

6.2.2: Neutron backgrounds

The possibility of neutrino-induced neutron backgrounds, presented in §5.2.4, was a somewhat unique concern for an SNS-based, low-threshold experiment. Section 5.5 discusses measurements carried out that provided a limit on the NINs production rate in a similar lead shielding structure as that used for the CsI[Na] CEνNS measurement.

Early work within The COHERENT Collaboration considered quantitatively the likely contribution from NIN events [77]. To address any concerns about NIN backgrounds, an additional layer of inner, hydrogen-rich shielding was added to the CsI[Na] shielding structure (§5.3). Subsequent MCNPX-PoliMi simulations confirmed that, with the additional shielding and with the production-rate limit developed using the *in situ* neutron measurements, NIN backgrounds were reduced to a negligible level [17].

Though the neutron measurements taken in place at the SNS suggest a meager contribution from prompt, SNS-beam-related neutrons, this background feature is included in the spectral models.

To determine the expected contribution to the spectra collected in the CsI[Na] detector during the CEνNS search, a neutron source described by the parameters determined in the prompt-neutron background measurement was simulated in the CEνNS-search geometry (§5.6). Using quenching factor and light yield information, the simulated distribution of nuclear-recoil energies was converted by P.S. Barbeau into observed photoelectrons. The real-valued count-rate distribution functions are shown in Fig. 6.3, normalized per gigawatt-hour of SNS run time.

6.2.3: CEνNS signals

The calculation of differential count-rate distributions for CEνNS events is discussed in Sec. 5.7, though the specific distribution used for subsequent analysis and in Ref. [17] was produced by P.S. Barbeau. Differential count rates in nuclear-recoil energy space, calculated using Eq. (1.1), are converted to photoelectron space using the photoelectron yield observed in the CsI[Na] detector deployed to the SNS (§5.1.1) and the CsI[Na] QF determined in Sec. 4.7 combining new and literature values. These distributions, produced individually for each of the three prominent neutrino components of the SNS flux (ν_μ, ν_e, $\overline{\nu}_\mu$), were passed through the acceptance cuts of both of the analysis pipelines to produce real-valued (i.e., not normalized to unity) PDFs for the count rate per photoelectron per gigawatt-hour of SNS exposure; the resulting distributions can be seen in Fig. 6.3.

The total expected count rate from CEνNS events, in aggregate and for each individual species, requires scaling of the PDFs shown in Fig. 6.3 by an appropriate exposure factor. These factors, expressed in gigawatt-hours, are unique between the two analyses and presented in Tab. 6.1. Much of the difference

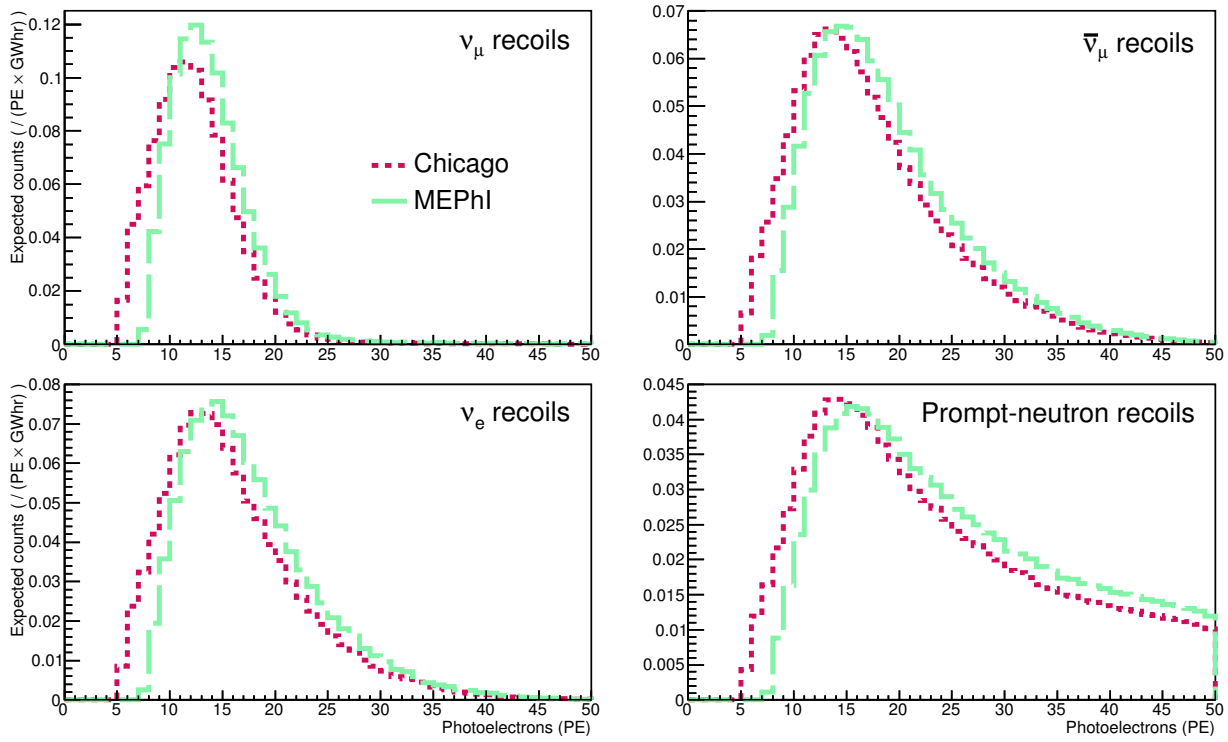

Figure 6.3: Expected observed photoelectron distributions, per gigawatt-hour of accumulated SNS run time, from CEνNS events of all flavors present in the SNS neutrino beam as well as recoils related to prompt SNS neutrons. For CEνNS signals, the Standard Model cross section is assumed. The light-green, coarsely-dashed lines show the distributions for the MEPhI/ITEP analysis and the crimson, finely-dashed lines show the same for the Chicago analysis; the corresponding acceptance curves can be seen in Figure 6.1. Distribution of these events in time is discussed in Section 6.3. The prompt neutrino population is composed entirely of the ν_μ signals while the delayed neutrino population is made up of both ν_e and $\overline{\nu}_\mu$ interactions.

Analysis	Beam energy exposure (GW·hr)
Chicago	7.48
MEPhI/ITEP	5.99

Table 6.1: Total beam energy exposure factors for the Chicago and MEPhI/ITEP analysis pipelines. These are used to scale the probability distribution functions which describe the CEνNS recoil spectra as determined by the Standard Model, shown in Figure 6.3. The differences in the exposure factors between the two analyses stem largely from the treatment of cut efficiencies. The determination of these two values is discussed more fully in the theses of B. Scholz [224] and A. Konovalov [156].

between the values of the two exposure factors originates in the treatment of live-times and cut efficiencies; these different treatments are also visible in the acceptance curves associated with the two pipelines, shown in Fig. 6.1. The determination of these exposure factors is closely related to the acceptance curves and is of fundamental importance to the overall analysis; this important work was carried out by B. Scholz and A. Konovalov, and detailed discussions are found in their respective theses, Ref. [224] and Ref. [156].

Section 6.3: Timing characteristics of SNS data

Each beam pulse from the SNS arrives with a distribution in time on the order of 100 ns; the timing features of the proton pulses are discussed in Sec. 2.4.1 and the proton bunch timing distribution averaged over the duration of the CsI[Na] CEνNS-search data collection can be seen in Fig. 2.1. Spallation, and thus production of the neutrons comprising the prompt neutron background, should very closely follow the proton pulse time distribution. Two distinct-in-time neutrino populations are expected: prompt neutrinos (ν_μ) associated with π^+ decay and delayed neutrinos (ν_e, $\overline{\nu}_\mu$) associated with muon decay (§2.5.1). The production of prompt neutrinos should relatively closely follow the proton distribution, having to account only for the 26-ns lifetime of the π^+. Muons, whose decay is described by a 2.2-μs exponential, are produced by the decay of the pions.

Following production, both the neutrinos and neutrons must travel from the spallation target to the CsI[Na] detector. Within the populations of both neutrons and neutrinos, the particles are assumed to propagate at the same velocity irrespective of energy. In the case of prompt, SNS neutrons, this assumption represents a scenario where the initial energies are sufficiently high to justify the energy-independent propagation time; only near the CsI[Na] system do these neutrons downscatter and take on the energy distribution used to model neutron backgrounds (§5.5.2).

Summarizing, the timing models for all components begin with the POT trace from Fig. 2.1. Describing first the *shapes* of the timing distributions, but not the absolute locations: the prompt neutron shape

follows the POT distribution exactly; the prompt neutrino distribution is the result of a convolution of the POT shape with a 26-ns exponential decay; and the delayed neutrino model is the *prompt* neutrino model convolved with a 2.2-μs exponential decay. These shapes are all then shifted by an appropriate value with respect to the DAQ trigger time; these shift values are based on simulation but backed up by the *in situ* neutron measurements described in Sec. 5.5.3. Figure 6.4 shows the timing distributions used for both neutrino populations, the prompt neutron component, and the original POT shape prior to shifting.

A final aspect of timing which is not included in the model used in this analysis is the spreading associated with the uncertain arrival time of the first-photoelectron in the waveform. The CEνNS data includes the integral number of photoelectrons and the 1st SPE arrival time, not simply "the interaction time". To explore the impact of this effect, a simple toy Monte Carlo was carried out for waveforms with n PEs. For each value of n between 1 and 20, 10000 sets of n "arrival times" were drawn randomly from an exponential model with time constant $\tau = 625$ ns, representing an approximation of characteristic timing for low-energy nuclear recoils [77]. Figure 6.5 shows the mean arrival time for the first PE from these simulations, as well as the time at which 95% of the photoelectrons have arrived; for signals with 5 photoelectrons, the point at which the acceptance efficiency becomes nonzero but is still very low, the first PE arrival time is within 100 ns. Combined with the coarse time binning (1 μs) adopted for the analysis, the arrival time statistics and the acceptance curves result in an associated uncertainty small compared to the dominant contributions of the QF and the neutrino flux; its impact is not included in this analysis.

Section 6.4: Analysis as a counting experiment

The most simple framework in which to analyze the CEνNS-search data is that of a counting experiment, where the data is reduced to scaler count data in the signal (coincident) and background (anti-coincident) regions. Comparison between the regions can be carried out to look for an excess of counts in the signal region. Such an analysis makes extremely limited use of the time and energy information that is recorded in the CsI[Na] experiment but provides an attractively simple approach to consideration of the results.

The only use of the time and energy information enters into the counting-experiment analysis in the selection of a region of interest (ROI). A very coarse ROI is inherently defined by the analysis pipelines, which restrict consideration to events with integrals between 0 and 50 PE occurring within 12 μs of the digitizer trigger; this ROI will be referred to as the "full range".

A cursory inspection of the distribution of CEνNS events, as predicted by the Standard Model, reveals that the full range includes regions of time and energy space which are likely to include few, if any, CEνNS counts; the SM-predicted CEνNS signal can be seen drawn as a heat map in Fig. 6.6. The significance of any

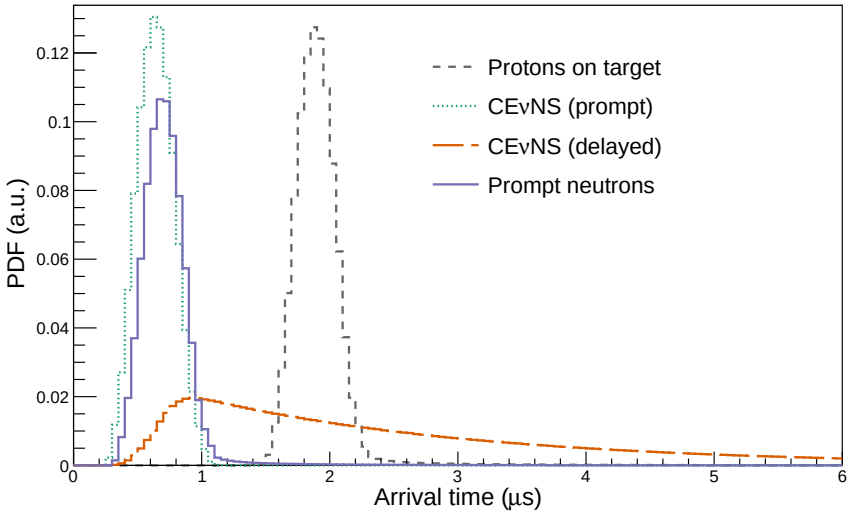

Figure 6.4: Timing distributions for neutrino interactions at the CsI[Na] detector at the SNS. The prompt neutrino component, drawn as a dotted, teal line, is associated with interactions from the ν_μ component of the neutrino population, which arises from the decay of stopped pions in the mercury target of the SNS; this time distribution is produced by convolving the protons-on-target PDF shown in Figure 2.1 with an exponential decay of $\tau = 26$ ns, the time constant for pion decay. Muons are produced in the pion decay, and the delayed neutrino population, shown by a dashed, burnt-orange line, is associated with stopped muon decay; this distribution is produced via convolution of the *prompt*-neutrino distribution with a 2.2-μs exponential. Neutrino production and associated timing are also discussed in Section 2.5 of the text.

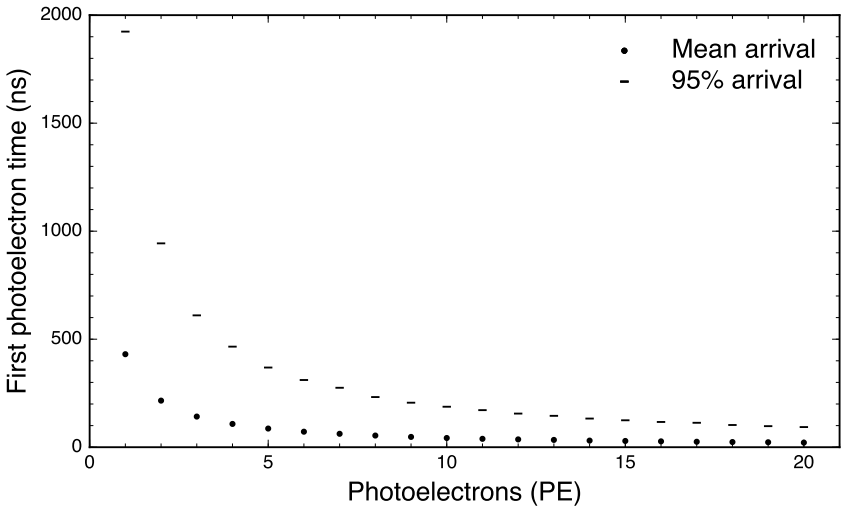

Figure 6.5: Time after interaction by which the first photoelectron of 50% and 95% of simulated CsI[Na] signals should have arrived. For signals with 5 PE, the smallest signals for which acceptance efficiency is nonzero in either analysis pipeline (see Section 6.1.2 and Figure 6.1), more than half of the signals should detect the first photoelectron within 100 ns. This data is generated by a toy Monte Carlo simulation of CsI[Na] signals modeled by a single exponential decay envelope with $\tau = 625$ ns, using an approximation from Collar *et al.* [77]. Considering the coarse 1-μs time binning used in this analysis, the arrival-time statistics combined with the acceptance efficiencies, and the substantially larger uncertainty associated with the quenching factor of CsI[Na] and the neutrino flux, the effects of arrival time are not included in the present results.

Analysis	Bin definition	Counts				Significance
		Signal	Background	Prompt neutrons	Residual	
Chicago	Full range	1207 ± 35	1032 ± 32	8.1 ± 2.0	167 ± 47	3.6σ
	ROI	547 ± 23	405 ± 20	6 ± 1.5	136 ± 31	4.4σ
MEPhI/ITEP	Full range	1210 ± 35	1078 ± 33	8.3 ± 2.1	124 ± 48	2.6σ
	ROI	538 ± 23	416 ± 20	5.7 ± 1.4	116 ± 31	3.8σ

Table 6.2: Results from both the Chicago and MEPhI/ITEP analyses of the CsI[Na] CEνNS search when analyzed as a single-bin counting experiment. Data using the full range of the data and the restricted ROI are represented. In the case of the contribution from prompt neutrons, the uncertainty represented here corresponds to the 25% uncertainty established in Section 5.5. Uncertainties on the signal and background counts are statistical.

signal excess could be improved by restricting the analysis region so that it includes less area with meager signal density. A rectangular region, spanning energy space between 6 and 30 PE and time space between 0 and 6 μs, is chosen and referred to as the restricted ROI.

Table 6.2 compiles the counting-experiment results from both the Chicago and MEPhI/ITEP analysis chains. Observation significance exceeds 4σ in the ROI of the Chicago analysis, having largely left unused any timing and energy information that the data may contain. Though the focus on the ROI makes some use of this information, the available information content exists at a much more finely grained level and can be utilized by higher-dimensional analysis.

Section 6.5: Analysis in one-dimensional space

A counting experiment only makes use of the scaler number of counts in a given ROI but there is more information available from the CsI[Na] experiment. Adding a single dimension to the analysis, distributions in energy space (with units of photoelectrons) and time space (with units of microseconds) can be exploited. As a 1-D analysis, these two bases are treated independently through the projection of the data and the models onto a single axis, e.g. the data and model are integrated in energy space to produce a time-space projection.

Data projected onto the PE axis can be seen in Fig. 6.7, along with spectral models; similar projections onto the arrival-time axis can be seen in Fig. 6.8.

A map of χ^2 values for the projected data can inform an understanding of both the best-fit number of counts and the preference for a CEνNS-like signal rather than its absence (i.e., the significance of the observation). For both analyses and both the time- and energy-space projections, χ^2 values are calculated using the full-range data and a variable number of CEνNS events; the χ^2 values as a function of CEνNS

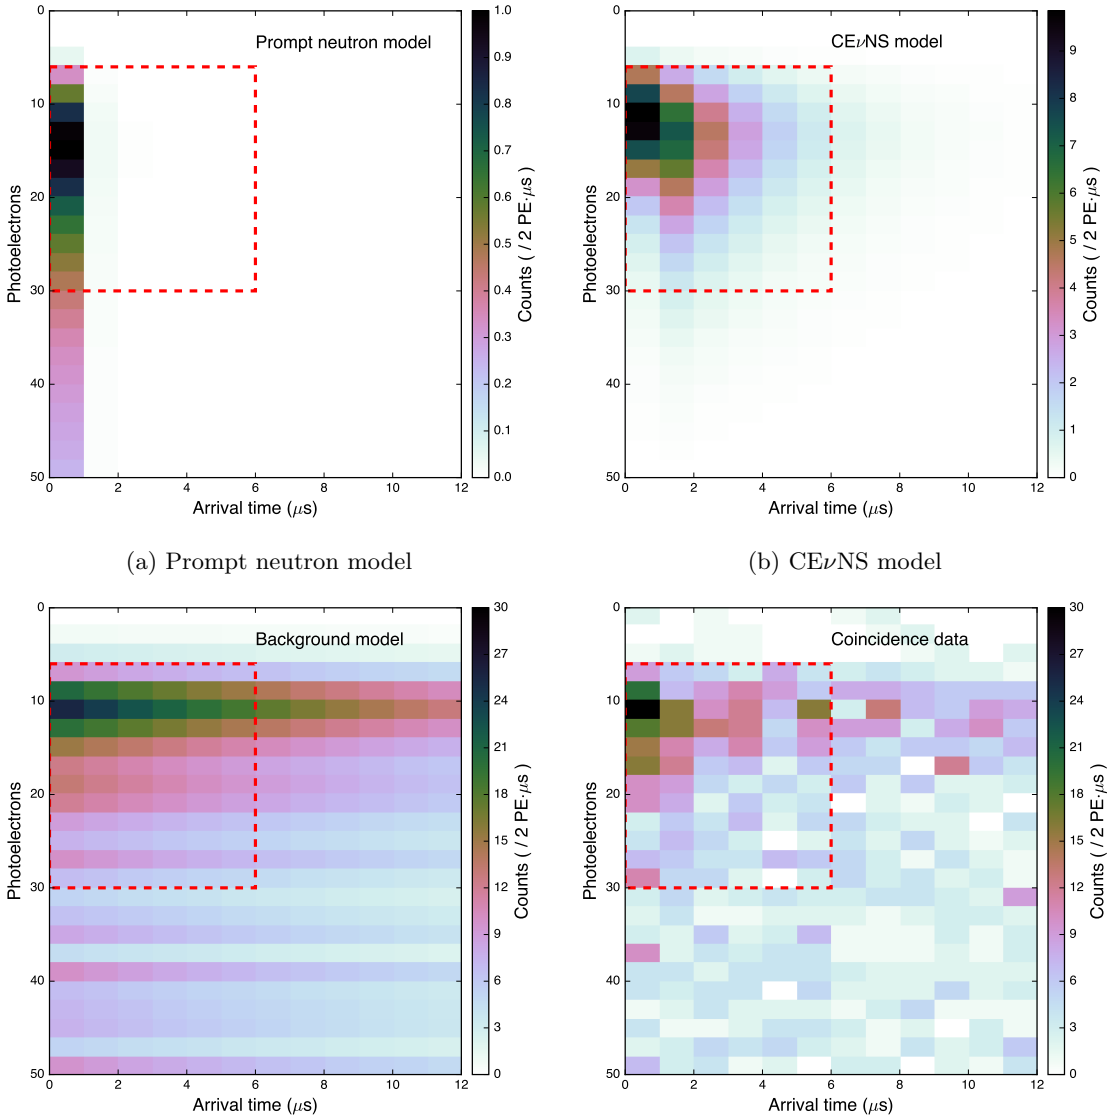

(a) Prompt neutron model

(b) CEνNS model

(c) Background model

(d) Coincidence region

Figure 6.6: Two-dimensional histograms showing (a) the prompt neutron model, (b) the combined CEνNS model consisting of each of the 3 constituent neutrino components, (c) the background model informed by the anticoincidence-region data (see Section 6.2.1), and (d) the coincidence-region data collected during the CEνNS search. The results in Table 6.2 reflect analysis over the entire range of the 2-D histograms ("Full range") and over the region indicated by the red dashed line ("ROI").

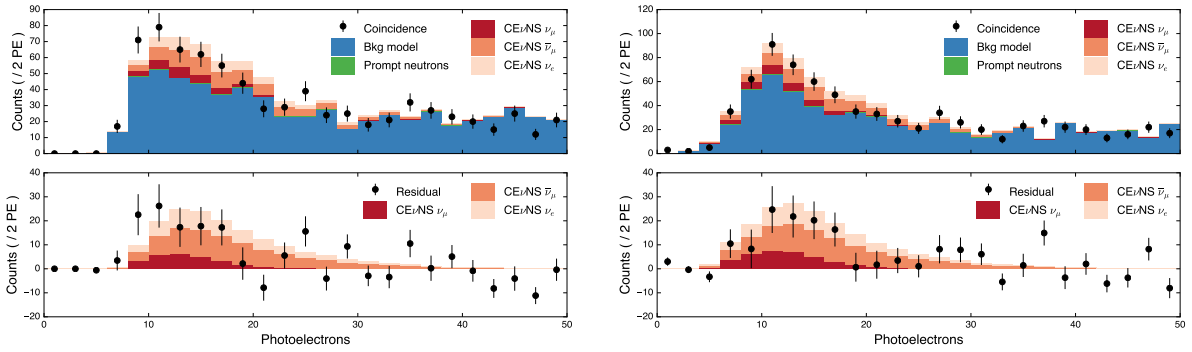

(a) MEPhI/ITEP analysis chain

(b) UChicago analysis chain

Figure 6.7: Photoelectron-space projection of CEνNS search data with best fit model overlaid. The lower panels show the residual data, with the AC data subtracted from the C data.

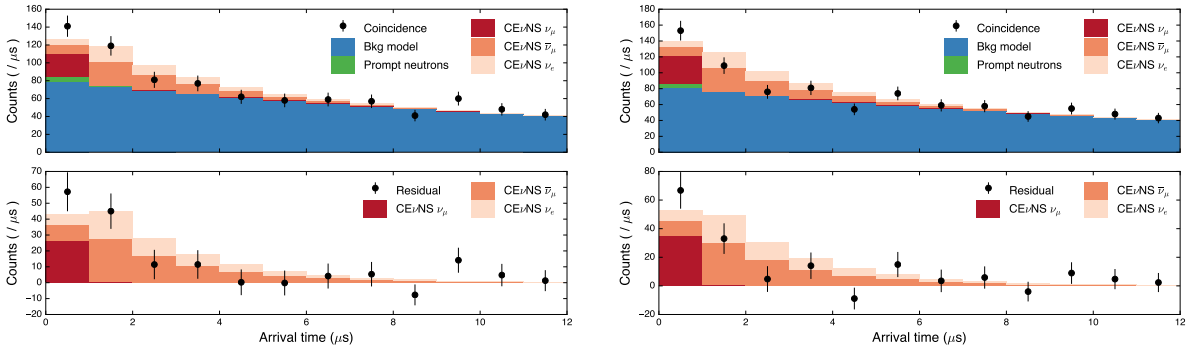

(a) MEPhI/ITEP analysis chain

(b) UChicago analysis chain

Figure 6.8: Arrival-time projection of CEνNS search data with best fit model overlaid. The lower panels show the residual data, with the AC data subtracted from the C data.

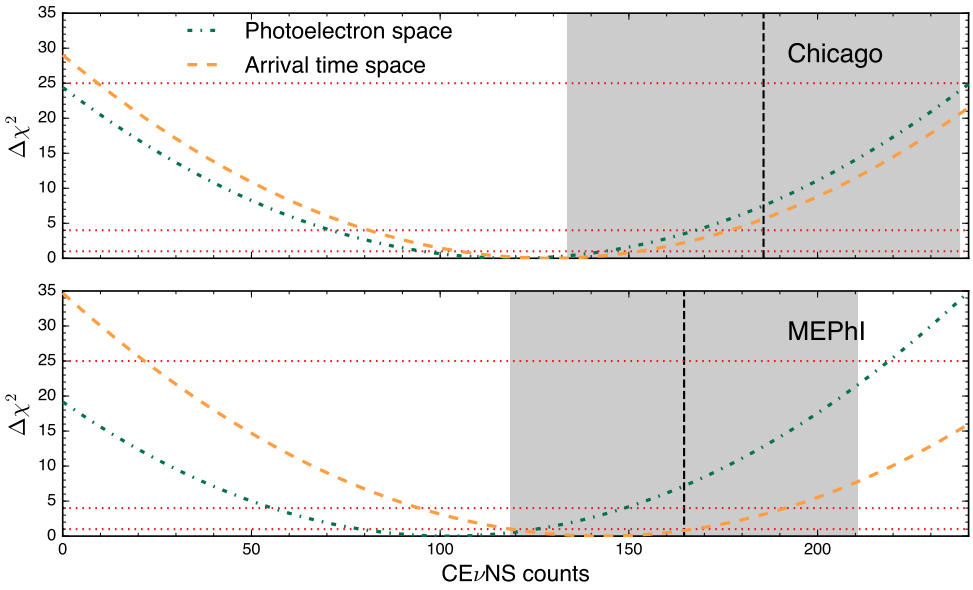

Figure 6.9: Results of χ^2 analysis for projections onto both the photoelectron and arrival-time axes; these are shown as green dash-dotted lines and dashed yellow. Results from the UChicago analysis are shown in the top panel, and results from the MEPhI/ITEP analysis are shown in the bottom panel. Standard Model predictions are shown as a gray vertical band; the vertical dashed black line indicated the mean SM prediction. Horizontal dotted red lines indicate the 68%, 95%, and 99.9999% confidence levels, corresponding to significance levels of 1σ, 3σ, and 5σ, around the best fit values.

counts can be seen in Fig. 6.9, where they are shifted so that the best-fit values are associated with $\chi^2 = 0$.

Table 6.3 collects the salient information from the χ^2 investigations for an observation of the CEνNS process. To discuss the observation significance from the χ^2 values, we define the quantity $\Delta\chi^2_0 \equiv \chi^2_0 - \chi^2_{\rm bf}$, where $\chi^2_{0({\rm bf})}$ is the value of χ^2 at 0 CEνNS counts (the number of CEνNS counts associated with the best fit); the χ^2 values drawn in Fig. 6.9 already perform a similar compensation, so values of $\Delta\chi^2_0$ for a given data appear as the χ^2 value at 0 CEνNS counts. The threshold for discovery in the particle physics community is generally considered "5σ": this corresponds to a situation where the observed effect (e.g., a collection of events in some energy range or at some time) can be explained by statistical fluctuations of other processes or backgrounds in *fewer* than 1 instance in ∼3.5 million [42]; the sigma value is described by the quantity $\sqrt{\Delta\chi^2_0}$.

For the time projections from both the Chicago and the MEPhI/ITEP analyses, the number of CEνNS events which best describe the data sets are associated with $\Delta\chi^2_0$ values meeting the criteria for declaring an observation of CEνNS. Still more powerful analyses can exploit the multidimensional nature of the data and signals and consider the numerous dimensions simultaneously, including any correlations.

Analysis	Dimension	Best fit	Standard Model	$\Delta\chi^2_0$	Pref. over null
Chicago	Energy	119 ± 24	186 ± 52	24	4.9σ
	Time	130 ± 24		29	5.4σ
MEPhI/ITEP	Energy	102 ± 23	165 ± 46	19	4.4σ
	Time	144 ± 24		35	5.9σ

Table 6.3: Results from χ^2 analysis of one-dimensional projections in time and energy space of CsI[Na] CEνNS search data. The threshold for observation or discovery is generally regarded to be 5σ, which is exceeded by the time-dimension projections for both analyses.

Section 6.6: Fit to two-dimensional data

6.6.1: Shape PDFs for fitting

Each of the components considered in the analysis (prompt neutrons, steady-state backgrounds, and CEνNS from each neutrino flavor) were represented in the fit to collected data by 2-dimensional PDFs describing their shapes in both time and photoelectron space. These features are discussed in Sections 6.2.3, 6.2.1, and 6.2.2. Figure 6.10 shows all of the 2-D PDFs used in the fit process: the prompt neutron model is in the upper-left corner; the aggregate CEνNS model is shown in the bottom middle position; and the steady-state background model is in the lower-right corner.

Each of these PDFs is additively combined to form the total model. The CEνNS PDFs were combined and treated as a single aggregate model; the rates of individual flavor contributions were not allowed vary independently, but the overall number of CEνNS counts[3] was allowed to float Overall normalization of the steady-state background PDF is constrained to be within statistical uncertainty of the number of counts present in the 2-D AC dataset. Finally, the number of prompt neutrons was constrained by the *in situ* neutron background measurements of Sec. 5.5, ultimately translating into a mean number of counts $\sim$ 1.1 counts/GW $\cdot$ hr $\pm 25\%$, in addition to Poisson errors, in the case of the Chicago analysis.

For the MEPhI/ITEP analysis, the same prescriptions are followed, but using the appropriate data to develop or constrain values.

6.6.2: Profile likelihood to address systematic uncertainties

In the two-dimensional analysis, taking full advantage of the energy and time distributions of the participating signals, we utilize a *binned maximum-likelihood* approach to fitting. Specifically, the fitting procedures

[3]The number of counts was defined as a fraction of the mean Standard Model estimate, but this was an implementation convenience and should not be of consequence.

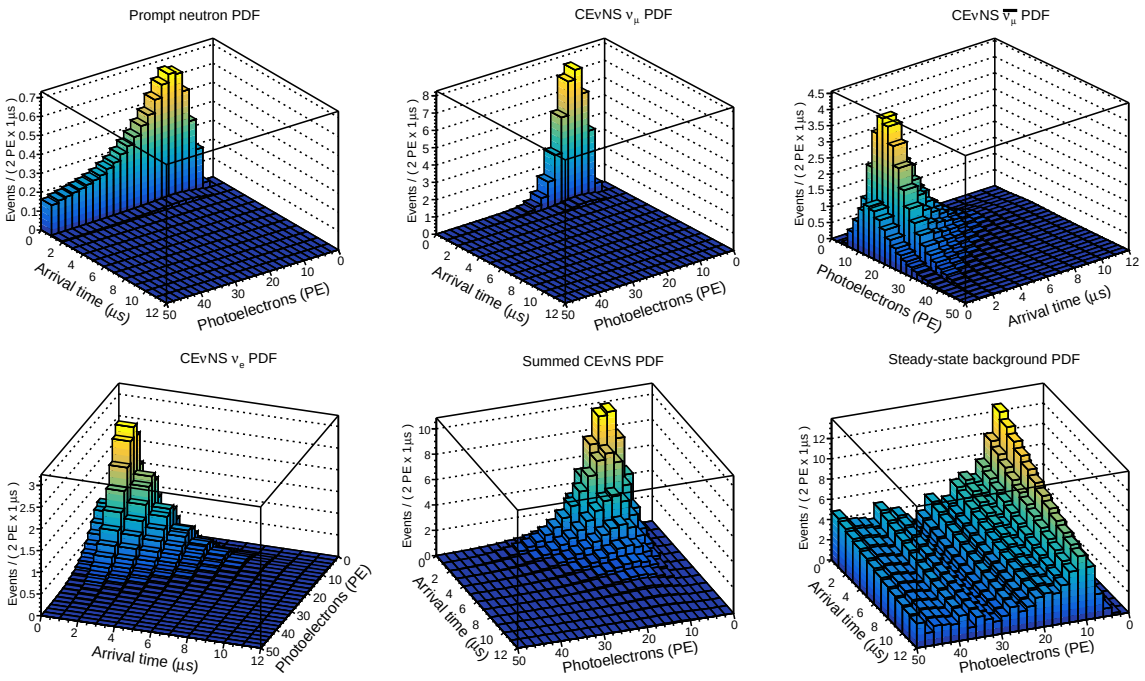

Figure 6.10: Two-dimensional probability distribution functions for the constituent features of the CEνNS search data from the 14.6-kg CsI[Na] detector at the SNS. These PDFs are summed together to produce the model which is fit to the collected data. These pertain to the UofC analysis pipeline; slightly different shapes, produced in the same way but using an alternative and appropriate acceptance efficiency curve, are used for the MEPhI/ITEP data. This figure appears in Reference [17].

used in this work *minimize* the *negative log likelihood* NLL: an equivalent procedure to maximizing the (non-negated) log likelihood which makes use of readily available numerical algorithms for locating minima. There are many excellent introductions to the maximum-likelihood formalism; a particularly apropos discussion can be found in sections 2.3 – 2.5 of Ref. [42]. An understanding of the mechanics of likelihood-based fitting is subsequently assumed.

The model used to fit the data is constructed in such a way that the parameter of interest, i.e. the one whose value gives us "our result", is the fraction of the Standard-Model-predicted CEνNS counts. Two additional parameters appear in the model: one associated with the number of prompt neutron events and the other associated with the number of steady-state background events. Speaking somewhat cavalierly, we may say that we "don't care" about these two additional parameters: the values they assume must adhere to whatever constraints we place upon them (directly or indirectly), but we are not seeking to produce a measurement of these parameters. In the context of our exploration of the number of CEνNS events present in the data, we would refer to these as *nuisance parameters*. Nuisance parameters appear in the accounting of systematic uncertainties; within the present model, they encode an inexact knowledge of the total number of both steady-state background and prompt neutron events[4].

Though not the approach utilized here, a discussion of nuisance parameters and their treatment should include some mention of marginalization [42, 233], which is the technique used in Bayesian analysis to handle this issue[5]. If the end goal of a Bayesian is to produce a distribution, say $f(X)$, which describes the parameter of interest X given some knowledge or evidence, then additional model parameters add additional dimensions to this final parameter space. A nuisance parameter is one that a Bayesian does not care to include in the final result: "I don't want to know X as a function of Y, I want to know X *given what I already know about* Y". Our intrepid Bayesian would seek to *marginalize over parameter* Y, using knowledge of Y given by some distribution $f(Y)$. Working in the framework of a multi-dimensional distribution, incorporating this knowledge of Y is conceptually straightforward: the Bayesian would simply integrate over Y.

Working in the frequentist statistical framework, we seek to *profile* over the nuisance parameters to incorporate their effects [42, 214]. Profiling consists of scanning over the parameter(s) of interest and, for each step in the scan, holding this parameter fixed while the NLL is minimized with all other parameters allowed to vary (subject to whatever, if any, respective constraint)[6]. The sequence of parameter-of-interest

[4]Shape uncertainties could also be included through the addition of additional nuisance parameters and associated reparameterization of the model; these modifications will be of importance in future analyses seeking to make statements beyond that of an observation of CEνNS (see a discussion in Sec. 7.2).

[5]Purists beware! The discussion of marginalization, and Bayesian statistics in general, is extremely informal and intended to provide a conceptual picture of the approach which may be more intuitive than that employed by frequentists.

[6]Practically speaking, one would generally first do an overall minimization of the NLL and profile over a limited range of parameter-of-interest values around the minimum.

values and the associated, minimized NLL values constitute the profile from which the best-fit value and uncertainty for the parameter of interest can be extracted.

Returning to a specific discussion of the CEνNS analysis, a profile likelihood analysis is carried out, using the described model and the beam-on, coincident data sample. No background subtraction is performed: the model incorporates the steady-state background, with its normalization among the model parameters. Profiling is performed over the parameter which defines the CEνNS-signal content of the model, defined in terms of the fraction of the Standard-Model prediction; the range of the scan is between SM fraction values which correspond to 0 and 300 CEνNS counts. The minimum NLL value in the profile curve defines the best-fit number of CEνNS counts. The uncertainty[7] on the best-fit value is defined by the locations at which the profile NLL has increased from the minimum value by 0.5 [214].

6.6.3: CEνNS search profile-likelihood results

Fits were carried out using data and models from both analysis pipelines (§6.1, §6.2). Figures 6.11 and 6.12 show the profile negative log likelihoods for the UofC analysis [224] and the MEPhI/ITEP analysis [156], respectively. In both cases, the best-fit values for observed CEνNS counts in the data, which are indicated by the locations at which the negative log-likelihood values are minimal, are within 1σ of the number of counts predicted by the Standard Model. The results of the profile likelihood analysis for both analysis pipelines are summarized in Tab. 6.4.

In Figs. 6.11 and 6.12 the NLL values are shifted by a single value, unique to each set of NLL values, so that the minima are shown with $-\ln(\mathcal{L}) = 0$. For visual clarity of the comparison presented in Fig. 6.13, each curve is shifted so that its respective minimum NLL value appears at 0.The different *absolute* NLL values for the two analyses are not informative. Conclusions and insight are gleaned from the relative trends within a single profile curve, especially the difference between the NLL values at the null hypothesis and the best-fit values of CEνNS counts, and this behavior is unaffected by the imposed shifts.

[7] On another practical note, the `minos` routine in `RooFit` will produce parameter uncertainties which correspond to those that would be found via profiling.

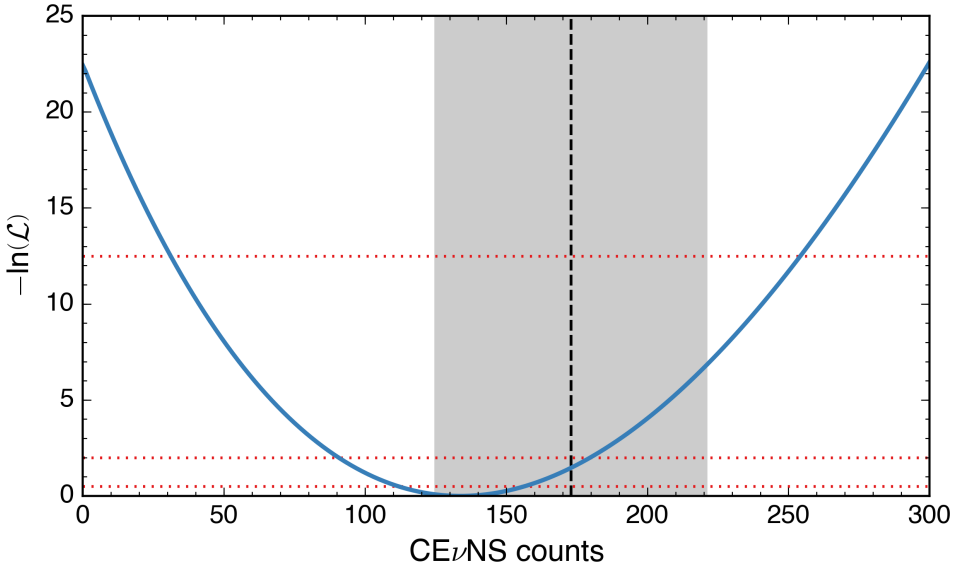

Figure 6.11: Negative profile log likelihood for the number of CEνNS events present in the COHERENT CsI[Na] data using the model described in the text. Likelihood values are shifted so that the best-fit value from the COHERENT data, 134 ± 22 CEνNS events, is drawn at 0. This result is within the 68% confidence band of the standard model prediction of 173 events, shown as a shaded region and a vertical, dashed line. The 68%, 95%, and 99.9999% confidence levels (1σ, 2σ, and 5σ) of the COHERENT fit are shown as ascending horizontal dotted lines. Comparison of log likelihood values at counts of 0 and 136 indicates that the null hypothesis, corresponding to an absence of CEνNS events, is rejected at a level of 6.7σ relative to the COHERENT best fit. A similar figure and caption appear in Reference [17].

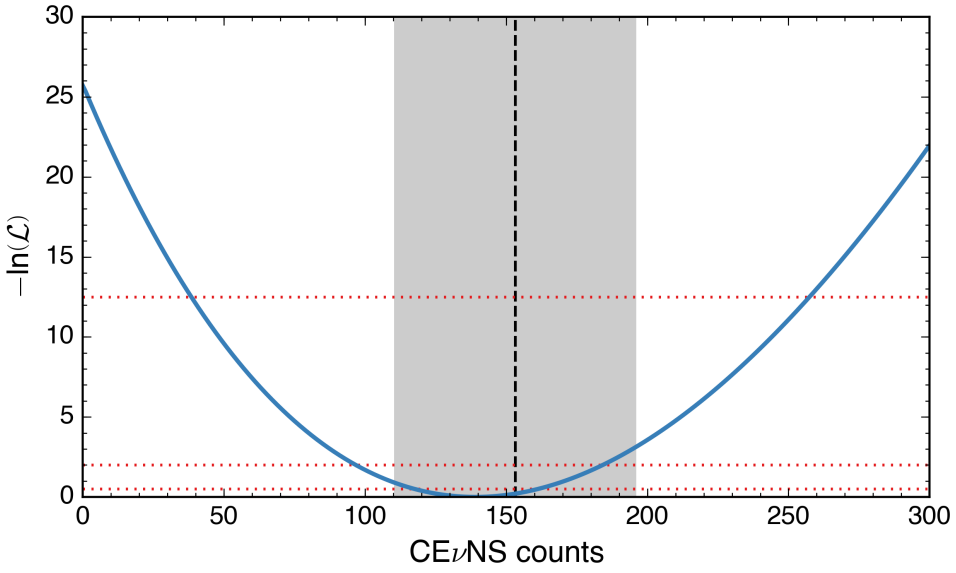

Figure 6.12: Like Figure 6.11 but for the analysis developed by A. Konovalov [156]. The 1σ confidence interval of the best-fit number of CEνNS counts, 132 ± 22, includes the Standard Model prediction of 153 ± 43. Preference for the best-fit value relative to the null hypothesis is 7.2σ.

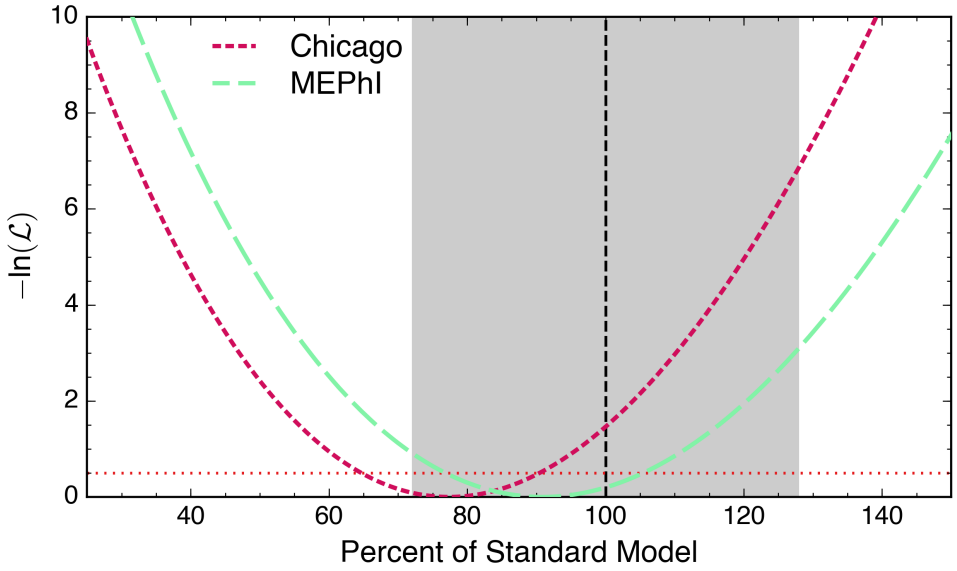

Figure 6.13: Comparison of profile likelihood results from the Chicago analysis, drawn as a finely-dashed crimson line, and the MEPhI/ITEP analysis, drawn as a coarsely-dashed light-green line. For the purpose of comparison, the horizontal axis is expressed in terms of the fraction of CEνNS counts predicted by the Standard Model for the respective analysis, rather than as an absolute number of counts. A dotted, horizontal, red line indicates the 1σ confidence levels for the best fit values for each analysis.

Analysis	SM prediction	Best fit	Best-fit preference over null
Chicago	173 ± 48	134 ± 22	6.7σ
MEPhI/ITEP	153 ± 43	132 ± 22	7.2σ

Table 6.4: Summary of the results of profile likelihood analyses carried out on both the Chicago and MEPhI/ITEP CsI[Na] CEνNS search analyses. The best-fit number of CEνNS counts is within 1σ of the Standard Model prediction.

CHAPTER 7: Additional CEνNS efforts and future opportunities

Section 7.1: Other CEνNS observation efforts

At the time of publication (Fall 2017), there are numerous proposed and ongoing efforts to perform an observation of the CEνNS process. Each of these experiments represent exceptional planning and effort, and the discussion presented here is limited only due to the indirect relevance to the work presented in this thesis.

Many of the contemporary efforts are based at nuclear reactors and thus rely on neutrino flavor and energy distributions unique from those of COHERENT or other experiments based at stopped-pion neutrino sources (§2.1.2). The MINER Collaboration [6] plans to use detectors developed as a part of the CDMS dark-matter detection effort, including silicon and germanium targets. Based at a unique research reactor facility, the features of the reactor and the utilization of multiple targets presents opportunities for suppression of some systematic challenges [6, 97]. Ricochet [167] is another experiment seeking to use CDMS-style detectors, but ultimately seeks to utilize much lower energy thresholds through alternative operation of these systems. Also proposed to take place at a research reactor, the very-low threshold of Ricochet could enable exciting physics searches (see, e.g., §§1.4, 1.5) that would be inaccessible to less-sensitive CEνNS measurements [167].

The CONNIE [104] and CONUS [62] Experiments are proposed to take place at commercial power reactors, which conceivably present substantially increased neutrino fluxes, though the realized gains may be offset by longer standoff distances and beset by other logistical or geometrical limitations. CONUS is based on another technology that has been utilized in dark-matter searches: germanium PPCs, such as those used in CoGeNT [1]. CONNIE utilizes thick silicon CCDs as its detection mechanism and could realize thresholds as low as $\sim$30 eV (ionizing energy), possibly producing a CEνNS measurement with small target masses $\mathcal{O}$(100 g).

COHERENT intends to perform additional measurements of the CEνNS process using different target nuclei and detection technologies. Each of the experiments referenced above rely on distinct technologies or implementations, and only germanium PPCs are represented in both external efforts and current plans for COHERENT; thus, measurements by all groups would afford tremendous diversity in the systematic uncertainties of each experiment. Alternative neutrino sources, such as reactors, present opportunities that warrant vigorous exploration by the experimental community. Additionally, as reproducibility is fundamental

to science, independent observation and measurement of the CEνNS process is essential.

Section 7.2: COHERENT beyond an initial observation of CEνNS

The work presented here and in Refs. [17, 224][1] represent the discovery of a CEνNS-like process but a critical requirement of an entirely unambiguous observation of the Standard Model CEνNS process still remains unfulfilled: observation of the N^2-like behavior of the cross section (§1.2). A component of the COHERENT experiment has always been a measurement of the CEνNS cross section with different target nuclei (§2.6), so despite the significant results thus far achieved there is still work to be done within the COHERENT effort.

Beyond the most basic goal of an unambiguous observation of CEνNS, many of the physics opportunities described in Chap. 1 require not necessarily measurements with different targets but much more precise measurements. Already with the present result the CEνNS cross section appears to be in agreement with the Standard Model value and meaningful constraints can be placed on neutrino-quark interactions indicative of physics beyond the Standard Model [17]. Any information regarding possible sterile neutrinos would require considerably greater statistics and could benefit from either a second detector or, less preferable, a second measurement at a different distance relative to the spallation target.The nuclear structure measurements described in Sec. 1.4 and the neutrino magnetic moment searches of Sec. 1.5.2 require analysis of recoil-energy spectra that will be inaccessible to the CsI[Na] detector; the exceptional low-noise characteristics of PPC-type germanium detectors would enable such experiments.

In the remaining sections, some improvements to the present analysis are proposed in addition to some expansions (both planned and already executed) of the experimental program of the COHERENT Collaboration which could yield valuable scientific returns.

Section 7.3: Improved CsI[Na] results

7.3.1: Refined understanding of CsI[Na] signal characteristics

The dominant source of uncertainty for the present CEνNS result arises from the determination of the CsI[Na] quenching factor, presenting a natural target for efforts aimed at increasing the precision of a subsequent result. Reanalysis of the COHERENT determinations of the QF (see Fig. 4.23) is warranted.Data collected in the experiment described in Chap. 4, with considerable digitization-length beyond the time at

[1] The analysis efforts already carried out by Alexey Konovalov, yielding "the MEPhI analysis" which has been discussed, will be supplemented and built upon in Ref. [156]; to avoid any speculation about the content of Ref. [156] beyond what is subject of discussion in this work, its citation is omitted in this section. This is not meant to diminish the contribution of A. Konovalov's efforts to the results presented.

which interactions in the CsI[Na] occur, present an information-rich data set for exploration on the effect of integration times on the QF.

Timing characteristics of signals from CsI[Na] are also an area where further investigation could pay dividends. Ultimately, an improved measurement of the timing distributions for low-energy nuclear recoils in CsI[Na] will be an important part of high-precision determinations of the CEνNS cross section (see "expanded parameter space" discussion in following section). Though not significant in the present analysis due to the very large QF uncertainty and smaller-but-still-large neutrino flux uncertainty, more robust models of CEνNS recoil spectra should account for variations due to arrival-time delays; for an example of the possible effects, see Fig. 7.5.

7.3.2: Higher-precision measurement of CEνNS cross section and comparison with Standard Model

Following the collection of the data which has informed the results in this work, the CsI[Na] detector has continued to acquire CEνNS-search data at the SNS: collection continues at the completion of writing (Fall 2017) with no discontinuation date planned. Increased statistics will invariably benefit the precision of the realizable result, but other techniques, some of which are detailed below, could be implemented in subsequent analyses that may yield more robust interpretations of the physics contained within the CEνNS cross section.

Alternative treatment of statistical effects in background model

Though the "factorized"-construction approach used here to develop a model for the steady-state background (see Sec. 6.2.1) is sound, a conceptually more straight-forward approach would be to use the 2-D AC data set as the background model itself while implementing appropriate treatment for bin-by-bin statistical fluctuations. Accounting for these fluctuations is the subject of a work by Barlow and Beeston [35] whose method (commonly referred to as "Beeston-Barlow", despite authorship ordering) is simple: each bin of a data set is treated as its own independent counting experiment, and individual-bin fluctuations are factored into the likelihood used during fitting. Despite the clarity of approach, implementation of the Beeston-Barlow method is computationally demanding.

With increased statistics, the need for the Beeston-Barlow treatment, or the factorized-construction treatment, is diminished. However, there will still be many bins (especially those at later arrival times or higher photoelectron numbers) where count numbers will be at a level that is ill-served by Gaussian approximations to statistical effects. The conceptual clarity of future analyses may be well served by implementation of the Beeston-Barlow method, allowing a simple statement of the origin of the background model ("the AC data

was used to represent the steady-state background, accounting for bin-by-bin statistical fluctuations using the method of Ref. [35]") rather than requiring a description of the factorization approach.

Unbinned likelihood analysis

Chi-square and (binned) maximum-likelihood analyses, though generally powerful and computationally fast, are subject to problems associated with the presence of zero-content bins. A more robust approach is offered by an *unbinned* likelihood analysis.

Unbinned likelihoods escape the problems associated with low statistics and can generally "make more for less (data)". The power of these techniques arises from the fact that each individual data point and its compatibility with the model parameters is considered in the calculation of the likelihood. By contrast, binned likelihoods group data points together (in bins) and then compare how many counts are expected by the model in a given bin with the content measured. Though both approaches employ "likelihood" functions, they are conceptually quite distinct. Since the unbinned approach requires an evaluation of the model at each data point, calculation of the likelihood function can be expensive[2], but with many free parameters and/or limited statistics, its effective use of *all* of the data, not just contents in bins, could pay dividends.

Expanded parameter space during fitting

The recoil rate distributions predicted for CEνNS interactions by Eq. (1.1) show, in some cases, both rate *and* shape differences depending on the values of various parameters. Ultimately, the observed distributions in terms of photoelectrons are subject also to possible shape changes due to effects such as quenching. The present analysis has allowed only for overall rate changes, after investigation of the shape-related effects suggested little significance.

A more robust model of the CEνNS signals would allow for variation of parameter values to result in alterations of the expected signal shapes. Additionally, "simultaneous" variation of multiple parameters will be important when attempting to infer any implications on, for instance, $\sin^2\theta_W$ in the presence of uncertainty on the nuclear form factor. Access to the wide array of physics accessible through CEνNS, manifested in the the differential cross section, also requires careful accounting of correlations, covariances, and any possible external constraints.

Computational requirements could become comparatively significant to those of the present analysis, which allowed for no shape dynamics. This is especially true for an unbinned approach, but both will be necessary to make meaningful progress into much of the physics portfolio described in Chap. 1.

[2]Consider, for example, a data set of 1M events: divided into 1000 bins and treated with a binned likelihood, each evaluation of the likelihood function requires 1000 evaluations of the PDF (1 and each bin center); the unbinned approach, however, must evaluate the PDF at each of the 1M data points.

Section 7.4: Candidate detector technologies for future COHERENT CEνNS measurements

Though the CsI[Na]-based search results presented here are consistent with the standard-model predicted CEνNS process, observation of the N^2 scaling of the cross section would provide additional certainty that the excess of observed events are attributable to CEνNS (see §2.6). COHERENT has included observation of the N^2 dependence as a priority and plans to accomplish this by making measurements of the CEνNS process at the SNS using several distinct detection systems which are constituted of target nuclei with differing N [14].

The plan at the time of writing, Fall 2017, is robust.The CENNS-10 detector is a 10-kg fiducial volume liquid argon detector that has been deployed to the SNS and is being advanced by COHERENT collaborators from Indiana University [58]. Another noble element detector, xenon-based RED-100 [15], is unlikely to see inclusion in the COHERENT experiment due to geopolitical tensions, despite earlier prospects for involvement [14]. Numerous members of the collaboration are forming plans to deploy ∼10 kg of germanium p-type point contact (PPC) detectors, a technology with a long connection to CEνNS but many uses in both the dark matter and $0\nu\beta\beta$ worlds. Finally, the lightest mass target will be in the form of the sodium component of NaI[Tl]; already, 185-kg of NaI[Tl] crystals are deployed to the SNS and collecting data focused on measurement of charged-current interactions on ^{127}I. The small array will serve as an R&D platform to develop CEνNS sensitivity, preceding the deployment of 2T of the crystals for a full-scale CEνNS measurement.

To roughly compare the recoil distributions expected from CEνNS events for these detectors, Fig. 7.1 shows the expected recoil count rate (per $\mathrm{kg}\cdot\mathrm{GW}\cdot\mathrm{hr}$) for the future detection targets. Rates in Fig. 7.1 are calculated with the Helm form factor and assume all detectors witness the same flux as the CsI[Na] detector (as if they were located at 19.3 m).

Section 7.5: Targets for future quenching-factor measurements at TUNL

Each of the detector technologies discussed in Sec. 7.4 as possible additions to the COHERENT suite of CEνNS experiments present targets of interest for independent measurements of nuclear-recoil quenching factors. In many cases, there is potential overlap with interest from the dark matter community: this is certainly true in the case of NaI[Tl], used in the DAMA/LIBRA experiment [45]; LAr, used in the DarkSide experiment [5]; and PPC germanium detectors, used in CoGeNT and its successors [1]. Additionally, among the original detectors planned for COHERENT was a dual-phase xenon TPC, RED-100 [16]; though geopolitical realities have changed these particular plans, the use of xenon in CEνNS searches is still attractive, and it is employed by several dark matter searches (i.e., LUX [13] and generations of the XENON experiment [22]), so a careful QF measurement is well motivated.

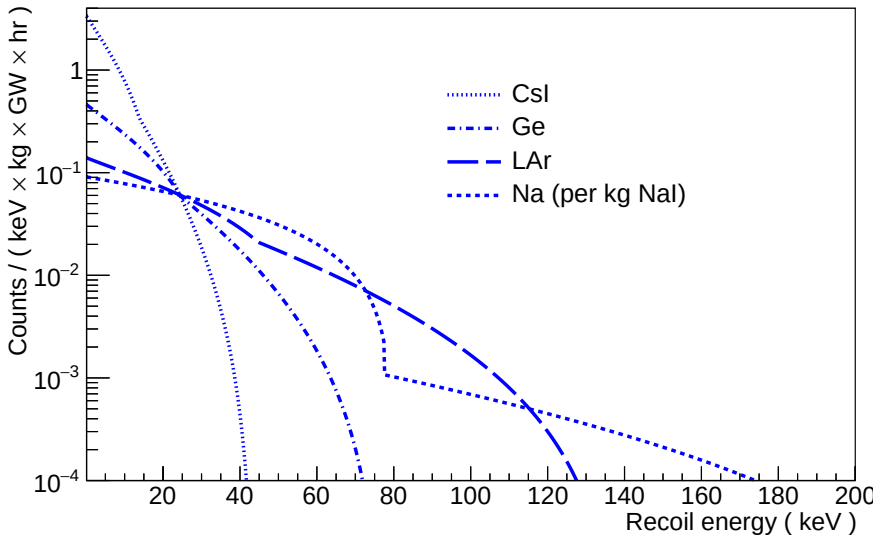

Figure 7.1: CEνNS recoil rate distributions for future detectors in the COHERENT suite. Rates are normalized per gigawatt-hour exposure at the SNS, assuming 0.08 ν/proton. For these estimates, each detector is treated as thought it witnesses the same flux as that incident upon the CsI[Na] detector. Lighter nuclear targets, such as Na or Ar, produce recoils at higher nuclear-recoil energies. These distributions can help understand the recoil-energy ROI for CEνNS and the sensitivity to uncertainty in the QF (see Section 3.2).

7.5.1: NaI[Tl]

An understanding of the response of NaI[Tl] to low-energy nuclear recoils is of considerable interest to the dark matter community due to its use in the DAMA/LIBRA WIMP search [45]. Considerable effort has been dedicated in recent years to developing such an understanding, including new measurements by Collar [75] and other groups [235, 257].

Data for such a measurement has already been collected at TUNL using a neutron-scattering configuration similar to that employed here to measure the QF for CsI[Na]. The ^{7}Li(p,n) reaction was used to produce a $\sim$580 keV neutron beam which was scattered into a circular array of 2"-diameter plastic scintillator detectors; these detectors do not posses PSD capabilities like the EJ-309 cells used in Chap. 4, but the relatively low energy neutrons may not have been easily distinguishable from γ-ray interactions in any detector. Preliminary data from this experiment is seen in Fig. 7.2 alongside literature values.

The TUNL measurement shown in Fig. 7.2 could be reanalyzed using the machinery employed in Chap. 4 in the analysis of CsI[Na] quenching factors. However, when this NaI[Tl] data was taken, an aluminum-backed LiF neutron-production target was used. Subsequent designs have used tantalum disks, rather than aluminum, and yield fewer beam-related γ-rays. The relatively high γ-ray flux from the aluminum target may impact the determination of the QF in ways that are not accurately reflected in the error bars shown in Fig. 7.2.

Figure 7.3 depicts the sodium recoil energy for different neutron energy and scattering angle configu-

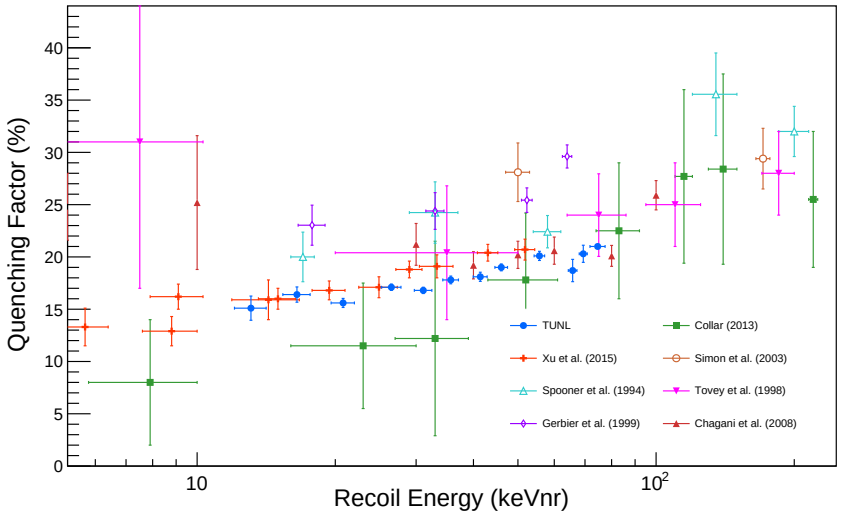

Figure 7.2: Values for the quenching factor of sodium recoils in NaI[Tl] from References [69, 75, 118, 232, 234, 240, 257]. Also included is the preliminary evaluation of data collected at TUNL using a setup similar to used in Chapter 4 with a ∼580 keV neutron beam produced with the $^{7}\mathrm{Li}(p,n)$ reaction. Techniques developed for the analysis of CsI[Na] QF data could be used to reevaluate this data.

rations. These relationships can be of assistance in planning a potential remeasurement of the QF for Na recoils in NaI[Tl] using a tantalum-backed LiF film for neutron production. The neutron beam energies appearing in Fig. 7.3 are generally accessible using the $^{7}\mathrm{Li}(p,n)$ neutron source.

7.5.2: Germanium and xenon

High-purity germanium detectors are used widely for their attractive spectroscopic capabilities [154]. P-type point contact germanium detectors, which offer enhanced capabilities relative to the common coaxial crystal configuration [34, 154], have found use in the challenging spaces of both WIMP dark-matter searches [1, 120] and $0\nu\beta\beta$ experiments [7, 198], and have thus been shown to have characteristics appropriate for a CEνNS search.

As a detection medium used by numerous WIMP searches, an understanding of the response by xenon to nuclear recoils has been of considerable interest to numerous groups. Several measurements exist [24, 176], and controversy has erupted [76] among the astroparticle physics community on the subject of proper analysis of QF measurements and their implications.

A general exploration of the accessible recoil energies in both Ge and Xe using an incident neutron beam is assisted by the kinematic plots of Fig. 7.4; these show contours of nuclear-recoil energy as a function of both incident-neutron energy and neutron-scattering angle.

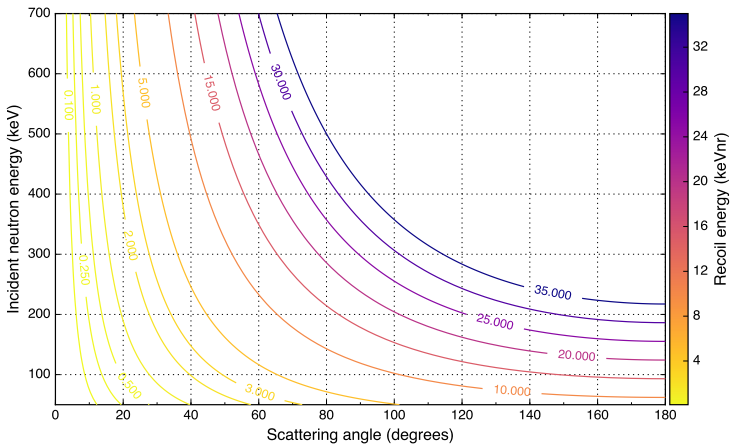

(a) Sodium recoil energies tagged as a function of recoil angle over a wide range of incident neutron energies.

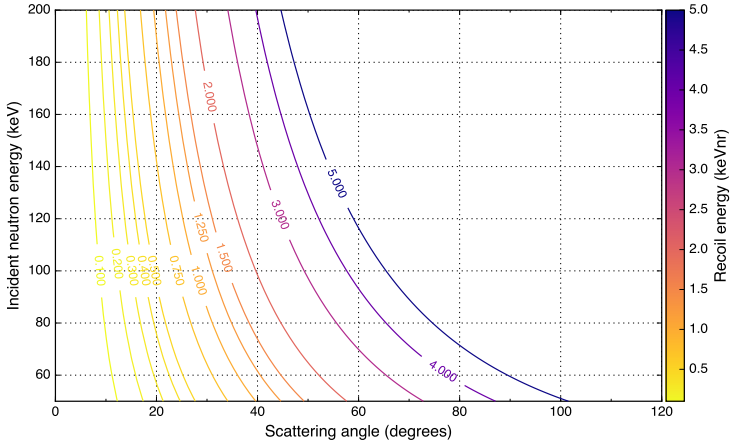

(b) Sodium recoil energies tagged as a function of recoil angle over a small range of low incident-neutron energies, highlighting recoil energies relevant to CEνNS.

Figure 7.3: Kinematic plots for possible quenching factor measurements in NaI[Tl], showing the sodium recoil energies as contours in the 2-dimensional space of incident-neutron energy and neutron-recoil angle. Plot (a) shows this for a wide range of incident-neutron energies while figure (b) highlights lower neutron energies and nuclear recoil energies more relevant to CEνNS signals.

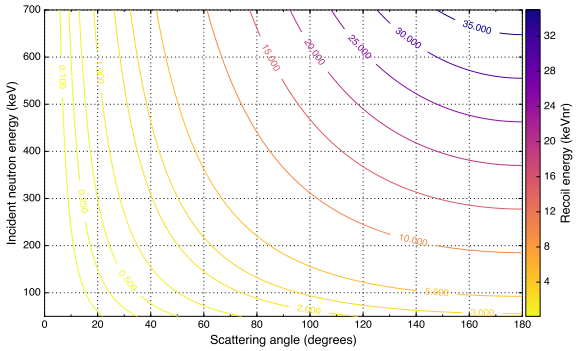

(a) Germanium recoil energies tagged as a function of recoil angle over a wide range of incident neutron energies.

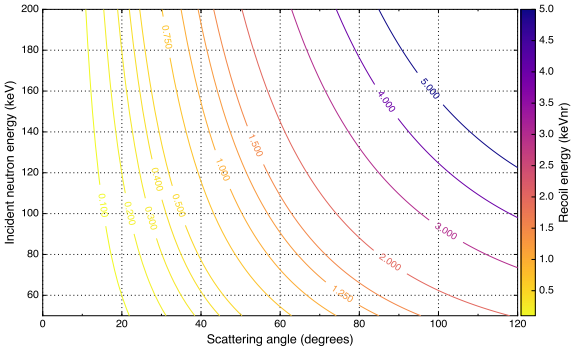

(b) Germanium recoil energies tagged as a function of recoil angle over a small range of low incident-neutron energies, highlighting recoil energies relevant to CEνNS.

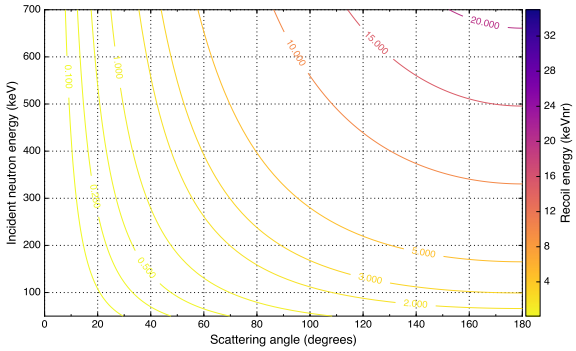

(c) Xenon recoil energies tagged as a function of recoil angle over a wide range of incident neutron energies.

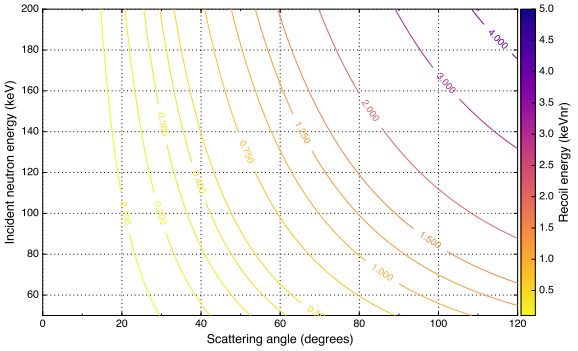

(d) Xenon recoil energies tagged as a function of recoil angle over a small range of low incident-neutron energies, highlighting recoil energies relevant to CEνNS.

Figure 7.4: Kinematic plots for possible quenching factor measurements in (a,b) germanium and (c,d) xenon, showing the recoil energies as contours in the 2-dimensional space of incident-neutron energy and neutron-recoil angle. Plots (a) and (c) show a wide range of incident-neutron energies while (b) and (d) highlight lower neutron energies and nuclear recoil energies more relevant to CEνNS signals.

7.5.3: General-use organic scintillators

In addition to explorations motivated by the astroparticle-physics community, the discussion in Sec. 4.4.2 illuminates a potentially valuable measurement: characterization of a fundamental, widely-generalizable quenching factor for EJ-309 liquid scintillator. Existing measurements show disagreement between efforts that use different size detectors: the results from Enqvist *et al.* are self-inconsistent with response measurements in 3 different-sized, cylindrical detectors, though it is possible that work was *intended* to report detector-specific response functions [103]. Tomanin *et al.*, using a cubic detector with sides of ∼10 cm, found a light yield inconsistent with the results of Enqvist *et al.* [239]. References [31, 144, 200] evaluated EJ-309 response in 2"-by-2" cylindrical cells, with [200] and [31] results in agreement and [144] data trending towards significantly higher yields at lower energies than other measurements. The geometry dependence on the existing determinations suggests that the reported quantities are not intrinsic to the scintillator, but also include geometric effects. Use of these measurements, then, by other experimenters is difficult as they would be unable to confidently deconvolve the detector-dependent response function involved in the production of the, ideally, intrinsic scintillator response.

Another way in which the existing EJ-309 response data could be improved is by providing experimental input for low ($<\sim$ 500 keV) neutrons. Tomanin *et al.* [239] do measure the response of their cell to 300 keV incident neutrons[3], but given the variation in existing data, a single data point at energies below ∼500 keV warrants supplement. The significance of low-energy data points is underscored by Pino *et al.*, who point out that choice of light-yield curve can quite directly affect the level at which a threshold is thought to be set[4] and this can have significant impact on total neutron detection efficiencies for $\mathcal{O}$(1 MeV) neutrons [200]. The problem presented by the dearth of sub-500-keV neutron data and the observation of Ref. [200] regarding thresholds is compounded by the study of Ref. [190] that highlights, in particular, divergence of extrapolated light-yield predictions at low energies depending on the choice of yield model.

There is significant variety in the experimental approaches that are used in previous characterization efforts. Many use time-of-flight to select specific neutron energies from broad spectra produced by radioactive sources (Refs. [31, 200]), accelerator-produced neutron beams (Ref. [239]), or both (Ref. [103]). Iwanowska *et al.* used tagged scattering to attempt to directly measure the response to individual proton recoils [144]. Other efforts rely on unfolding procedures to infer the yield from recoils (the "response" to 4 MeV neutrons, for instance, is made up of many proton, and possibly carbon, recoils), but do not directly observe individual

[3]In both the abstract and the conclusion of their paper, the authors of Ref. [239] suggest a measurement down to 500 keV neutron energy; in the body of the text, however, they report 300 keV. The reported minimum bombarding-proton energy of 2.05 MeV, used with the ^{7}Li(p,n) neutron-production reaction, should correspond to $\sim$ 291 keV neutrons according to Ref. [172].

[4]This is not unlike the threshold-related effect that QFs have on CEνNS measurements (§3.2).

recoil events; this procedure can be challenging [190], and may introduce model dependencies. The light yields from Ref. [144] were produced using relatively high-energy neutrons (∼14 MeV), and the authors identify this as the source of significant uncertainty at the lowest proton-recoil energies. In light of the potentially "cleaner" measurement of proton-recoil light yields offered by tagged scattering, the tension between the results of Ref. [144] and other efforts warrants further investigation.

A measurement which produces a detailed description of the fundamental ratio between electron-recoil and proton/carbon-recoil events for EJ-309 could be of great utility to experimenters in both fundamental and applied physics research, as well as members of the nuclear security and engineering communities. With these responses, experimenters would need only to understand the geometric effects of their particular detectors, which could then be convolved with the reported light yield functions to produce reliable nuclear-recoil response models. Given the scatter in literature values, it would be important to perform measurements with different size detectors and confirm that the determined response function generalizes effectively. The neutron beam facility at TUNL is well equipped to produce such a measurement, using a range of incident neutron energies.

Both the endpoint and tagged-scattering approaches of QF measurements (§3.3.3) could be used to build confidence and cross check values; additionally, many different experimental schemes could be imagined to increase efficiency of utilized beam time. For instance, multiple detectors could be placed at different angles away from 0° if an uncollimated neutron beam is used, allowing simultaneous endpoint response measurements. A similar effect could be accomplished with the use of a collimated neutron beam, like the one used in the CsI[Na] QF measurement described in this work, if a central "radiator" is placed in the beam path; multiple EJ-309 detectors could be placed at angles around the radiator, witnessing the scattered beam neutrons and subject to fewer beam-in-room-related backgrounds (e.g., beam neutrons which enter the experimental area, scatter and moderate in the concrete walls or floor, and submerge the detectors in a bath of thermal neutrons). If the radiator is a fast, scintillating detector itself, time-of-flight between the radiator and the backing detectors (under characterization, in this case) could be used to determine the neutron energy incident upon the backing detector on an event-by-event basis, thus allowing the use of a broad spectrum (and high flux) neutron beam. In another possible efficiency multiplier, the radiating target itself could be an EJ-309 cell, and tagged-scatter QF measurements could be extracted from its signals, not simply a timing input for analysis of backing-detector data.

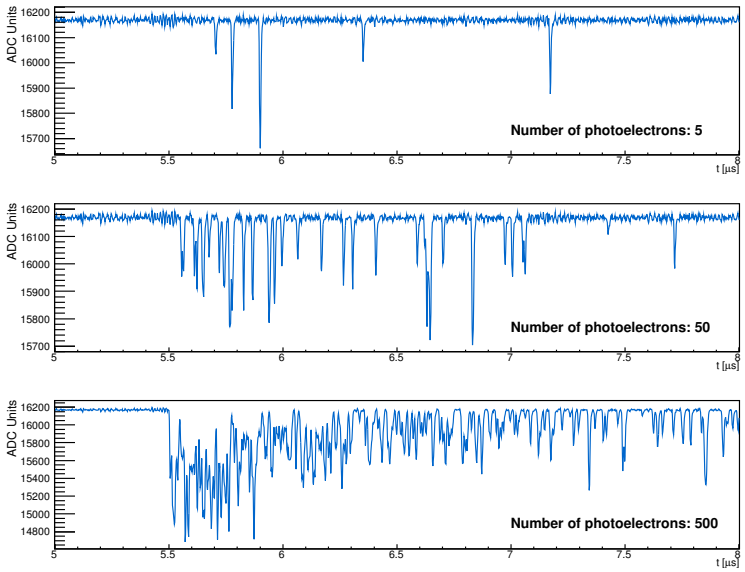

Figure 7.5: Simulated waveforms with different numbers of total photoelectrons demonstrating the significant variation in the observed timing characteristics of signals in the CsI[Na] detector. These are generated with the Monte Carlo approach described in Section C.5, assuming scintillation light in the detector is governed by two exponentials with time constants $\tau_{\text{fast}} = 589 \pm 4$ ns and $\tau_{\text{slow}} = 6.7 \pm 2.4$ μs and a signal intensity ratio between the slow and fast distributions of $I_{\text{slow}}/I_{\text{fast}} \sim 0.41$, as reported previously by Collar *et al.* [77]. In each waveform, the "interaction" occurs at the same time (5500 ns), but the first photoelectron observed in the waveforms arrives at some probabilistically-determined time after this. The possibility of the first photoelectron occurring appreciably after the interaction for few-photoelectron signals is visible in the top panel.

Section 7.6: Improvements to experimental approach for quenching factors and measurement of directional effects

7.6.1: Unbiased signal region identification

Many previous efforts at measuring quenching factors have employed some event-selection criteria involving the detector under characterization. For example, Ref. [75] employed a coincidence-window requirement between the first photoelectron observed in the scatterer and the tagged-neutron event in the backing detector to identify a signal-candidate events. Such a requirement introduces a bias in hardware on the allowed time for arrival of photoelectrons in the scatterer PMT. For energy depositions of a sufficient level, such a requirement and its associated bias are of little concern: at high photoelectron number statistics, the observed signal from a PMT will begin to resemble the timing envelope that defines whatever scintillating material is under evaluation. However, for very-low energy depositions, where the entire observed signal may consist of $<\sim$10 photoelectrons, the requirement on observation of any of these few PEs is no longer innocuous. Simulated waveforms illustrating this phenomenon are shown in Fig 7.5.

Both the COHERENT QF results for CsI[Na], presented in Chap. 4 and also in Ref. [224], were triggered

entirely based on information from the backing detectors, which ensures that minimal hardware bias is introduced at the time of acquisition. The present analyses, results of which are shown in Fig. 4.23, require identification of PE signals in the CsI[Na] waveform during post-processing. Data collected as described in Chap. 4 can, however, identify using entirely external information where the signal region should be in the CsI[Na] waveform. This is enabled thanks to the pulsed beam operation and the digitization of the BPM signal. Though earlier efforts have attempted to carefully account for any bias introduced by triggering configurations [75], analysis using this externally defined signal region should represent a QF measurement with absolutely minimal triggering bias and would be a valuable contribution to the literature, especially in light of its unique systematic independence.

7.6.2: Search for channeling

In the past, crystalline organic scintillators, such as stilbene mentioned as a scintillator with potential utility in a broad application space (§7.5.3), have been shown to feature anisotropic light yields [53]: the signal per unit energy deposition depends on certain geometric factors. For use in low-energy neutron detection, consideration of this effect could lead experimenters to orient detector crystals in such as way as to maximize the signal that is expected. Conversely, failure to consider the anisotropy could lead to misinterpretation of collected signals.

Potential for misinterpretation due to anisotropic yields extends to experiments such as a CEνNS search or attempts to detect dark matter. It has been speculated that the DAMA/LIBRA dark matter signal could be result of "channeling", a directional effect where a recoiling nucleus experiences reduced stopping power as a consequence of traveling through a "channel" in the crystal lattice of the detector [75].

Attempts to observe the existence of channeling effects in NaI[Tl] have thus far shown null results [75]. Using crystalline organic scintillators to demonstrate and qualify the capability to measure directional yield effects, the facility for quenching factor measurements at TUNL presents a testbed at which the absence of channeling in NaI[Tl] can be confirmed with confidence. Directional effects in dark matter detectors are considered a way of gaining more confidence in any WIMP detection, as they may enable "pointing" to a source and possibly avoid the neutrino floor discussed in Sec. 1.6 [191], so identification of anisotropies in inorganic crystalline scintillators could be valuable.

Section 7.7: Exploration of thermal spike explanation for quenching

As discussed in Sec. 4.7, present theoretical models for scintillation light production in inorganic scintillators do not predict the trend towards lower quenching factors at lower nuclear recoil energies that have

been observed in numerous crystals [242].Ionization quenching, an effect where high ionization density in a crystal results in annhiliation of exciton-hole pairs, is an effect that is considered in the Birks model for scintillation [53] which informs, for example, the model of Tretyak [242].

It is possible that the subtleties associated with the thermal spike process described in Sec. 3.1.1 are not entirely accounted for by simply considering regions of particularly high ionization density. Modern molecular dynamics studies focused on WIMP detection in bubble chambers have found some results that differ from the picture presented by Seitz regarding the thermal spike model [89]. Additionally, thermal spikes have been shown to result in localized, transient phase changes [181]. If a region of the crystal, localized around the interaction site, were to have some transient change of phase or alteration of the lattice structure, this could result in changes to exciton dynamics that are not described well by a treatment which is grounded in ionization quenching. An exploration of this would likely require full molecular dynamics simulations to understand in good detail the material response of the crystal. Perhaps critical recoil energies, corresponding to local temperatures, could be identified where changes might be expected , and those regions could be explored with great care.

Section 7.8: Measurement of the neutrino-induced neutron process at the SNS

Neutrino-induced neutron (NIN) spallation is a background for the COHERENT collaboration discussed in §5.2. To facilitate the CsI[Na]-based CEνNS measurement presented here, it has been sufficient to establish that the NIN contribution is negligible with the salient detector and shielding geometry (§5.5), but the process of neutrino-induced neutron production is of considerably broader interest than simply as a background for other neutrino experiments at the SNS.

7.8.1: Underlying physics and implications of NINs

Neutrons can be produced by both charged- and neutral-current interactions on heavy nuclei such as lead. Examples of neutron-producing reactions are [100]

$$\begin{aligned}
\nu_e + {}^{208}\mathrm{Pb} &\rightarrow {}^{208}\mathrm{Bi}^* + e^- \\
{}^{208}\mathrm{Bi}^* &\rightarrow {}^{208-y}\mathrm{Bi} + x\gamma + yn, \\
\nu_x + {}^{208}\mathrm{Pb} &\rightarrow {}^{208}\mathrm{Pb}^* + \nu_x' \\
{}^{208}\mathrm{Pb}^* &\rightarrow {}^{208-y}\mathrm{Pb} + x\gamma + yn,
\end{aligned}$$

where the first sequence demonstrates a charged-current route and the second a neutral-current example. Theoretical calculations of this process have considerable uncertainties, on the order of a factor of 3 [102, 114, 155]. Indeed, even calculations of the inclusive charged-current cross sections for nuclei show substantial disagreement between different nuclear models [193]. This provides ample opportunity for contribution from any measurement of the NIN process on a wide range of targets.

Astrophysical environments rich with both heavy elements and an appreciable neutrino flux exist and are often candidate locations for the rapid neutron capture process, r-process. As these locations provide the essential ingredients for a NINs process to take place, it is ultimately possible that NINs could play a role in nucleosynthesis of heavy elements in supernovae [207, 255]. Additionally, the NINS process has been implicated as a potential mechanism which, ultimately, through supernova neutrino detection, can elucidate the neutrino mass hierarchy [245].

Finally, the NINs process is the mechanism underlying the functionality of the Helium and Lead Observatory (HALO) supernovae neutrino detector [96, 244]. With the current state of theoretical uncertainty on the NIN cross section, some experimental constraint could establish greater confidence in the efficiency for detection of HALO. A measurement using stopped-pion neutrinos from the SNS will not be well matched with the expected energy distribution of supernova neutrinos [100], so care must be taken in extrapolating any established constraint.

7.8.2: NINs measurement at the SNS: the neutrino cubes

While the *in situ* neutron background measurement discussed in Sec. 5.5 could produce an upper limit on the NINs cross section from lead, a dedicated measurement was designed in the form of the lead "neutrino cube", or "nube". The lead nube is, at its core, a ∼1000-kg mass of lead with cavities into which neutron detectors can be inserted; see design schematics in Fig. 7.6.

Though the first fully instrumented nube was focused on measuring the NIN process on lead, the same geometry can be used to measure the NIN process on a wide range of targets, replacing the lead volume with another material. Some care needs to be exercised when making such substitutions however. During the initial design, simulations were carried out for both lead and iron nubes: the two material have considerably different neutron transport properties, so detection efficiencies can vary substantially.

The neutron detectors used were liquid scintillator cells custom ordered from Eljen Technology. Each detector has a cavity with a 4.5" diameter and length of 9" which is filled with Eljen EJ-301 liquid scintillator, comparable to Bicron BC-501A; the cells are outfitted with an Electron Tubes 9821-KEB 3"-PMT, which are well-suited for PSD applications [211]. Prior to installation in the nube, these detectors were subjected

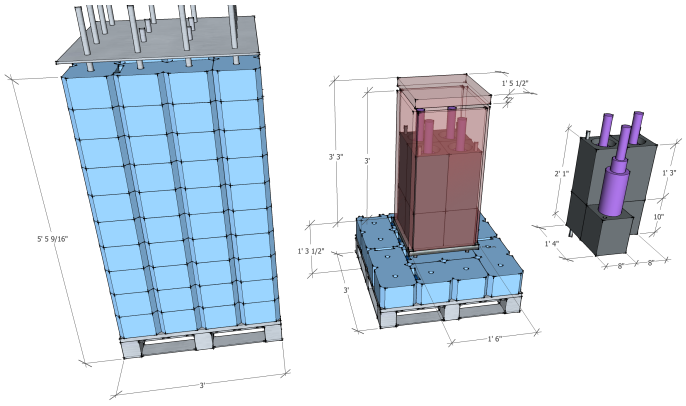

(a) Exploded view

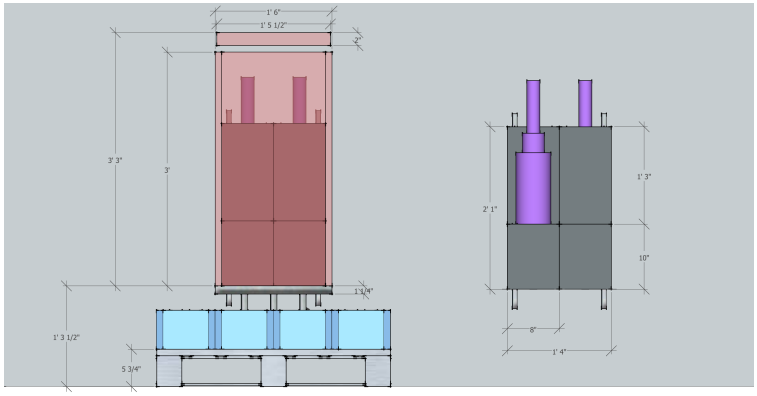

(b) Profile view of nube internals

Figure 7.6: Schematics showing the design of the lead neutrino cube. Muon-veto panels are shown translucent blush and simple models of the large Eljen liquid scintillator assemblies are rendered in amethyst. PMTs and holders for the muon-veto panels are not shown, along with the piece of foam supporting the top veto panel above the vertical panels [211].

Figure 7.7: Photograph taken looking along the proton beam direction just upstream of the neutron production target, which is obscured from view and held in place behind the visible flange in the bottom right of the image. The 4 large detectors visible in the middle-left of the image are the custom EJ-301 cells from Eljen. Quasi-monoenergetic neutron beams were produced using the $^{7}\mathrm{Li}(p,n)$ reaction; several different proton energies were used to produce neutron beams of different energies. As each detector was located at a slightly different neutron-emission angle with respect to the production target, each detector was subject to slightly different beam energies. From Reference [211].

to a characterization effort at TUNL. To form an understanding of the neutron PSD performance of these detectors, quasi-monoenergetic neutron beams were produced in the TUNL nTOF room (see Fig. 4.1, lower right) using a pulsed proton beam incident upon a LiF film deposited on an aluminum backing. A photograph from this characterization effort can be seen in Fig. 7.7.

Following characterization, the large LS cells were moved along with the remainder of the experiment to the SNS. Since September 2015, the lead nube has been collecting neutrino data. There have been preliminary efforts within the collaboration to analyze the data, but the first result from the neutrino cube experiments is to be developed. Work will be continued on this effort.

APPENDIX A: NEUTRON DETECTION WITH LIQUID SCINTILLATOR CELLS

Section A.1: The scintillation process in organic scintillators

Detection of both the unscattered and scattered neutrons is essential to quenching factor measurement described here and is accomplished with small liquid scintillator cells. Given its significance to the experiment, we discuss briefly the scintillation process in organic scintillators and the origin of pulse-shape discrimination capabilities in these materials. Detail is limited here; for a full discussion, refer to Birks [53].

A.1.1: General process of scintillation from organic compounds

The process of scintillation in organic scintillators is distinct from that of inorganic scintillators discussed in Chap. 3 where dopants are (generally) directly responsible for the production of light. Instead, in organic compounds, scintillation arises from de-excitation of π electrons, which participate in a particular kind of bond (π) in organic molecules [53].

Systems of π electrons have excitation levels associated with singlet ($S_0, S_1, S_2, \dots$) and triplet ($T_1, T_2, \dots$) states, each with vibrational sublevels [53]; level spacing between the ground and 1^{st} singlet states is generally between 3 and 4 eV, and the vibrational level spacing is typically on the order of 0.15 eV [154]. Direct excitations from the ground state, S_{00}, to the lowest-lying triplet state T_{10} are forbidden, so excitations are generally to excited singlet levels [53]. The fluorescence process of scintillation in organic compounds produces photons through transition from the 1^{st} singlet state to one of the vibrational states associated with the ground state configuration, e.g. $S_{10} \rightarrow S_{01}$; the transition from these low-lying vibrational states to the lowest-energy configuration will be non-radiative [53]. Organic scintillators are generally transparent to their own scintillation because the photons released when decaying to a vibrational state, e.g. $S_{10} \rightarrow S_{01}$, will be of insufficient energy to excite a molecule from the *ground* state S_{00} to the S_1 states [154].

In addition to the fluorescence process, the phosphorescence process[1] produces light via transition from an excited triplet state to the ground singlet state [53]. The T_1 state is much longer lived than S_1, so this process generally yields light over a much longer period of time [53]. It is not well understood how "intersystem crossing", the process by which singlet states transfer their energy to triplet states, proceeds, but evidence has long existed that such transitions occur [53].

[1] See Birks [53] for a discussion on the terminology.

A.1.2: Stopping-power dependence on fluorescence and phosphorescence: the origin of PSD

An important property of many organic scintillators, including those used in these measurements, is that ionizing particles with different stopping powers dE/dx show different balances of fluorescence and phosphorescence, or prompt and delayed scintillation light; specifically, particles with greater ionization densities (higher dE/dx) seem to have a larger fraction of delayed light relative to lower dE/dx particles [53]. This may be explained by interactions between excited and unexcited π electrons, which have an increased likelihood of occurring in a region of higher ionization density [154]. Whatever the origin of the effect, it results in the ability to perform pulse shape discrimination (PSD), where interactions involving particles of considerably different stopping powers are separated from one other in some parameter space, typically called the pulse shape parameter (PSP). In this work, PSD is used to distinguish electrons generated by γ-ray interactions (which have relatively low dE/dx values) from recoiling proton or carbon nuclei, with relatively high dE/dx, generated by neutron scattering.

A.1.3: Light yield in EJ-309 liquid scintillator

The signal yield characteristics of EJ-309, a general-use liquid scintillator with good PSD performance, have been well studied, but we focus here on the characterization performed by Pino *et al.* [200], who employ a parameterization for the light yield like that used in other studies of both EJ-309 [103] and the broader class of "NE213-equivalent" liquid scintillators [159]. In this treatment, the electron-equivalent light yield L_p for proton recoils in EJ-309 is given by

$$L_{ee}\left(E_p\right) = aE_p - b\left[1 - \exp\left(-cE_p^d\right)\right], \tag{A.1}$$

with

$$a = 0.62 \pm 0.03, \quad b = 1.3 \pm 0.3, \quad c = 0.39 \pm 0.08, \quad d = 0.97 \pm 0.04,$$

where E_p is the proton-recoil energy; and a, b, c, and d are fit parameters [200], though only the central values are used in the present analysis, neglecting the uncertainties.

This treatment is subject to some deficiencies. Enqvist *et al.* saw considerable variation as a function of detector geometry in their determination of EJ-309 light yield, attributing these variations to "reduced self-attenuation and less time spread for the shorter light paths before conversion to electric pulses" [103]. The values used here, from Ref. [200], are associated with a similarly sized EJ-309 cell: both those used in this experiment and that evaluated by Pino *et al.* are 2"-diameter, 2"-long right circular cylinders.

Section A.2: Realization of pulse shape discrimination in organic scintillators

A.2.1: The charge-integration approach

The different ratios of fluorescence to phosphorescence between electron and nuclear-recoil energy depositions can be exploited to discriminate between γ-ray and neutron interactions. To quantitatively discuss the PSD of a detector, we begin by subdividing each scintillator pulse into two regions in time: the "head" and the "tail", corresponding to windows intended to preferentially contain fluorescence and phosphorescence, respectively. The γ- or neutron-like qualities of a pulse can then be discussed in terms of a pulse shape parameter or PSD parameter,

$$\mathrm{PSP} \equiv \frac{\int_{t_{\mathrm{tail}}}^{t_f} Q(t)\,dt}{\int_{t_0}^{t_f} Q(t)\,dt}, \tag{A.2}$$

where $Q(t)$ is the charge measured from a PMT observing the scintillator at time t and t_0, t_{tail}, and t_f are the start time of the pulse, the start time of the tail region of the pulse, and the full length of the integration window of the pulse, respectively [212]. Such a quantification is referred to as the charge-integration approach to PSD [212].

Reliable performance of the charge-integration method depends on accurate determination of the pulse onset time, which serves as a reference against which t_0, t_{tail}, and t_f are all defined. Effects such as timing walk can negatively impact the discriminating power as this results in the tail integration region being shifted with respect to the pulse. Consequently, implementation of the charge-integration technique should be accompanied by a stable, walk-limiting timing circuit or algorithm, such as a constant-fraction discriminator.

A.2.2: The Mesytec MPD-4 PSD module

The MPD-4 is a single-width NIM module with 4 channels, each of which has a constant-fraction discriminator and performs firmware pulse-shape discrimination calculations with adjustable PSD window settings covering a range appropriate for liquid scintillators. Each channel of an MPD-4 has two analog voltage outputs, corresponding to the integrated charge of a signal and its pulse-shape parameter, in addition to a logic-level output appropriate for gating a peak-sensing ADC around the analog signals of the channel and a fast, selectable trigger output which can generate logic pulses whenever the PSP of the channel meets certain criteria (neutron-like, γ-like, or all). Additionally, the MPD-4 module has a common trigger output which represents the logical OR of the triggers from each channel in the module.

APPENDIX B: PHOTOELECTRON DISTRIBUTIONS

Section B.1: Probability distribution models for few-PE signals in PMTs

B.1.1: The Pólya distribution

To describe the charge distribution associated with single- and multi-photoelectron emission from the cathode of a PMT, we seek a distribution which is physically motivated by the underlying signal generation process in PMTs: amplification of the number of electrons, originating from one or more photoelectrons emitted by the photocathode, through successive striking of a number of dynodes. Prescott provides a viable model in the form of the Pólya distribution, which is described by [205]

$$P(x) = \frac{\mu^x}{x!} (1 + b\mu)^{-x-1/b} \prod_{i=1}^{x-1} (1 + ib), \tag{B.1}$$

including the generating function

$$G(s) = [1 + b\mu(1 - s)]^{-1/b}, \tag{B.2}$$

with parameters b and μ, the latter of which describes the mean. This distribution also arises in the context of wire chambers, where it is employed to describe the gain fluctuations in electron-avalanche amplification in the presence of nonuniform electric fields; addressing this application, Blum *et al.* [56] offer an alternative parametrization

$$P(n) = \frac{1}{\overline{n}} \frac{(\theta+1)^{\theta+1}}{\Gamma(\theta+1)} \left(\frac{n}{\overline{n}}\right)^{\theta} \exp\left[-(\theta+1)\frac{n}{\overline{n}}\right] \tag{B.3}$$

where

$$\sigma^2 = \frac{\overline{n}^2}{\theta+1}.$$

B.1.2: The Gamma distribution

The Pólya distribution addresses many of the needs for a model of SPE response: it is physically motivated, approaches suitable distributions at extremes (i.e., exponential or gaussian), and can be expressed in a reasonably simple analytical form. However, to be generally useful, spectra containing multiple photoelectrons must be described by additive combinations of nPE components, each of which is itself described by a convolution of n SPE distributions. Numerical or FFT-based convolutions of PDFs can be used to build an approximate PDF for multi-PE events, but using the resulting PDFs in the context of fitting a spectrum is computationally expensive.

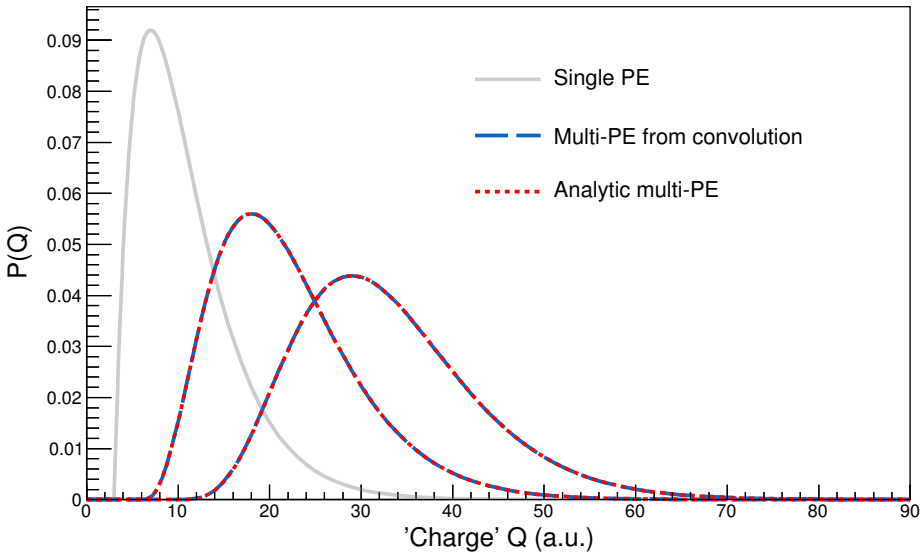

Figure B.1: Models for photoelectron shapes using the Gamma distribution in `RooFit`, `RooGamma`. Two- and three-photoelectron distributions are drawn, calculated both by convolution of SPEs and by analytical expression, the latter of which is more computationally efficient; agreement between these two calculation approaches is evident. The arbitrary parameter values supplied to the `RooGamma` function defining the SPE shape shown here are $\gamma = 2,\ \beta = 4,\ \mu = 3$.

	Single PE PDF	n PE PDF
Parameter	γ	$n \times \gamma$
	β	β
	μ	$n \times \mu$

Table B.1: Relationship between parameter values of an SPE distribution and an nPE distribution, when describing photoelectron PDFs using Gamma distributions. The parameters use the naming convention of the `RooGamma` distribution in `RooFit`.

A path forward is available by noting that the Pólya distribution is known by several alternative names: the compound Poisson distribution; the negative binomial distribution; and the Gamma distribution. Appropriately, the Gamma distribution has been found to describe few-PE spectra from PMTs [236] and used successfully to describe a similar amplification process present in electron-multiplying charge-coupled devices (EMCCDs), such as those used in observational astronomy [127].

Figure B.1 shows multi-photoelectron shapes produced by convolution of SPE shapes and by analytical expressions, showing agreement between the two descriptions. To describe nPE shapes in `RooFit` using a `RooGamma` distribution, the relationship illustrated in Tab. B.1 can be used, fully describing the parameters for a distribution with n PEs based on the SPE shape.

APPENDIX C: WAVEFORM FILTERING

Section C.1: Filtering of digitized waveforms

The processing of time-series data is a relevant concern across nearly every quantitative discipline, yielding a vast collection of application- and theoretically-minded literature. In any real scenario, an inevitable consideration must address the presence of noise in these data. A bestiary of filters or denoising techniques would reference a dizzying range of fields, and most filters will have particular strengths or weaknesses, better addressing different combinations of signal and noise components. The dynamic range of signals relevant to the analysis of the CsI[Na] QF measurements demands robust filter performance over a wide range of both frequency and amplitude space.

Signals from the lowest-energy elastic scattering events are expected to produce only a few photoelectrons ($\mathcal{O}(1)$) while the inelastic scattering events will be composed of $\mathcal{O}(1000)$ PEs; when combined with the individual PE timing characteristics and the PE arrival time distribution, the net result is that this range of signals is decidedly non-homogenous in any domain. Despite the differences in character, it is important that signals from both of these populations be included in the analysis, as the inelastic scattering events provide a useful normalization point for signal yield. Consequently, "tuning" of a single filter to provide consistent performance for the diverse signal populations relevant to this experiment must be done with care. Adding to this challenge is the fact that the SPE timing characteristics are not substantially dissimilar from representative timescales of the readily-observable, high-frequency baseline noise present in the experimental data; consequently, a time-domain filter targeting this noise may have a considerable impact on the measured SPE charge – an unwelcome feature.

Section C.2: Single photoelectron charge distributions and the need for filtering

The importance of filtering is demonstrated very well by examination of integral distributions from the CsI[Na] detector, focusing on the low-charge region where single- and few-photoelectron signals are prominent features. Given the low-energy signals of interest to these QF measurements and CEvNS searches, this is a region of great interest: the ability to determine the mean charge of a single photoelectron is of fundamental importance for analysis of measurements where the physics signals themselves are composed of only a handful of photoelectrons. This SPE charge distribution fitting procedure, discussed in detail in Section 4.3.3, is predicated on discernible features in the few-PE spectrum.

Figure C.1 shows the few-PE charge distributions for CsI[Na] waveforms subject to three different analysis processes: baseline-subtraction only, with no filtering, using a baseline determined by the normal mode

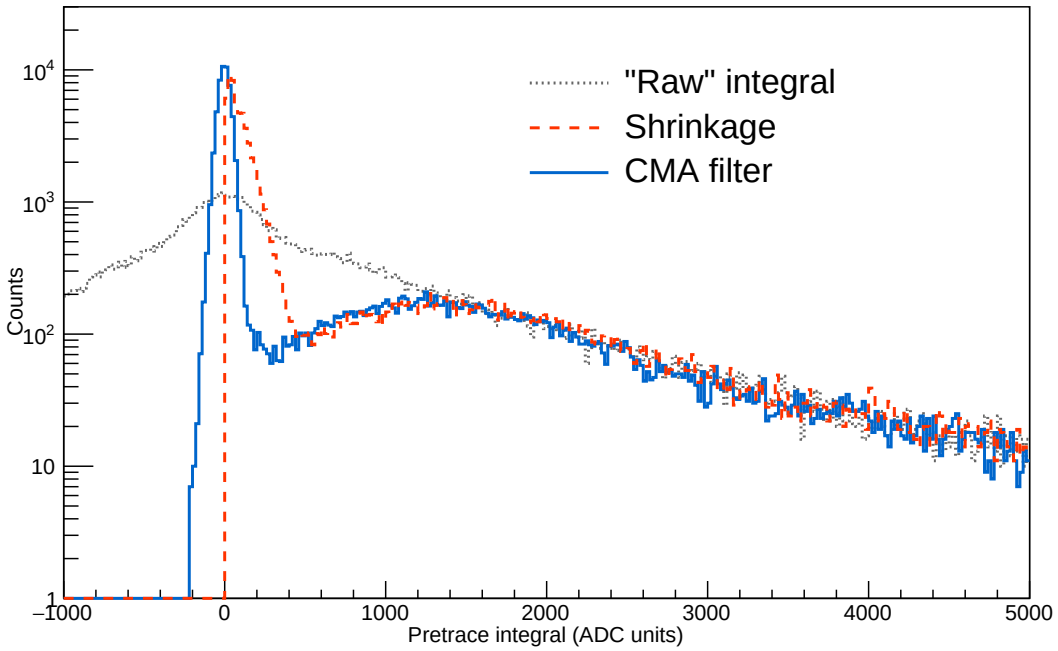

Figure C.1: Charge distributions in the few-photoelectron range from the CsI[Na] detector with several treatments of the waveforms. Shown as a dashed, gray line is the integral taken with a baseline determined by the "normal mode approximation" described in Section C.4.2. Application of the baseline shrinkage approach of Section C.3 results in the distribution shown as a dotted, red line. The CMA filtered waveforms produce the distribution seen as a azure line. The features of the CMA filtered distribution are thoroughly well described by physically motivated models while the shrinkage technique results in an obviously-truncated pedestal component, potentially betraying more subtle effects that could mislead a determination of the SPE charge. The charge distribution from the "unfiltered" waveforms shows significant degradation of spectral features and demonstrates the need for filtering if SPEs are to be reliably identified in integrated-signal space.

approximation of Sec. C.4.2; hard-threshold shrinkage, described in Sec. C.3; and CMA filtering presented in Sec. C.4. The single most prominent feature in these distributions is the pedestal: the peaked distribution centered, ideally, around 0, which results from integrating baseline in the absence of any signal. Comparison of the pedestal shapes of the three distributions highlights the important differences between the approaches and advantages of using the CMA filter. The loss of resolution suffered by the unfiltered approach is apparent. Though the "shrunk" and CMA distributions are largely indistinguishable at integrals above ∼2000 ADC units, the pedestal is more narrowly resolved by the CMA analysis and the pedestal is artificially truncated at ∼0 by the shrinkage approach. The truncation draws attention to the fact that the near-pedestal features are distorted by shrinkage, which reduces the applicable range of the cleanly-defined analytical model used to fit this region and determine SPE charge. Prioritization of the recovery of pedestal-region spectrum characteristics, mindful of the need to accurately determine the mean charge of SPEs, motivates adoption of the CMA filter approach.

Section C.3: Baseline shrinkage

A relatively straightforward approach to noise reduction employs an idea similar to the concept of *shrinkage* with hard thresholding from the realm of wavelet denoising [92]. To determine an integral using this approach, following determination of the baseline (the choice of algorithm is not of specific importance), an amplitude threshold relative to this baseline is chosen and any waveform sample with a value below threshold will *not* be included in the integration.

Such an approach has the benefit of being conceptually simple. However, when using a physically-motivated model for the expected SPE-level integral distributions, modifications must be made to account for the effects introduced by "shrinking" the waveform. The effect is most obviously visible on the pedestal component of the charge distribution, which is truncated at $\sim$0 in the shrinking process. Comparison with the CMA distribution in Fig. C.1 reveals that the characteristics of the remaining pedestal may also be affected by shrinkage.

Section C.4: The conditional moving average (CMA) filter

To address the needs presented by these experiments that proved difficult to address with other filtering approaches[1], the conditional moving average (CMA) filter was developed[2]. The CMA filter is a variation on the standard moving average, or "boxcar", filter that provides an effective and conceptually straightforward approach for smoothing a noisy waveform. Typically, a boxcar filter would be used as a low-pass filter, removing high-frequency fluctuations from a signal. The CMA filter described here works similarly, but it is also sharply biased against deviation from its "present" value at any given point along the waveform: this feature helps to avoid filtering out both high and, particularly, low frequency components of signal.

An important terminological clarification should be offered here: **the term "CMA filter" and variations thereof are used to refer to a two-step process and the final output of these processes**. These steps are:

1. **Baseline determination** The baseline of the waveform is determined by direct application of the CMA algorithm. This baseline can include complicated structure, but any abrupt changes will be rejected.

2. **Baseline subtraction (and inversion)** The determined baseline is subtracted from the original waveform. If relevant, the resulting value is multiplied by -1, so that originally negative-going signals

[1] An entirely exhaustive exploration of the signal processing literature was not conducted, but many of the "most obvious" filtering algorithms did not show promise.

[2] "CMA" is also used to refer to a cumulative moving average filter.

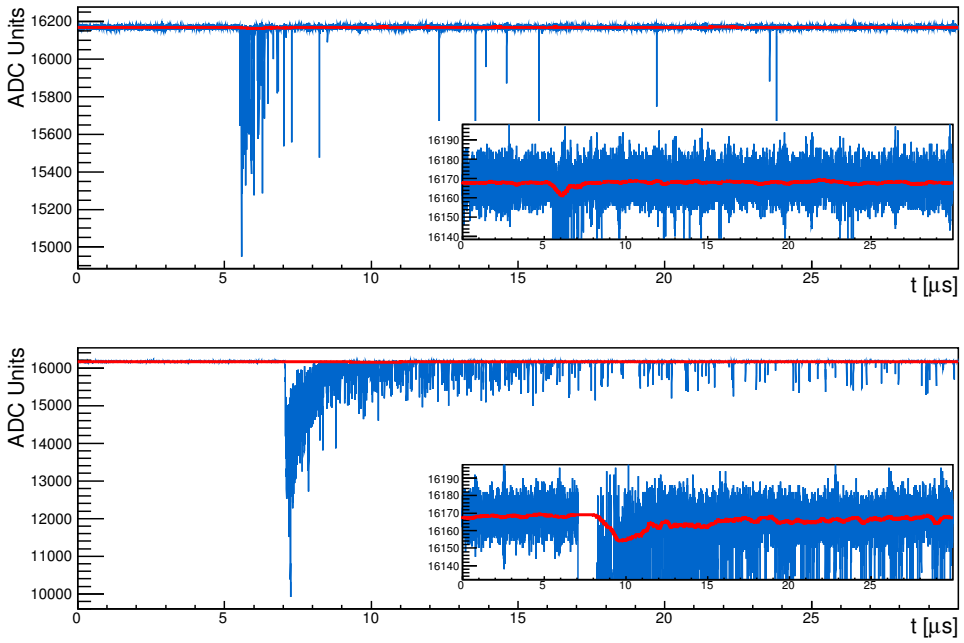

Figure C.2: Real waveforms associated with the scattering crystal from the CsI[Na] experiment for events with substantially different energy depositions. Also shown, in red, is the baseline as determined by the conditional moving average (CMA) filter described in Section C.4 of the text. The insets show a zoomed region around the baselines, better demonstrating the response of the CMA filter to the presence of signal. The filter shows little response to the presence of isolated single photoelectron pulses, as can be seen at later times in the waveform shown at top, corresponding to a lower-energy signal.

are made positive going.

It is more natural in the present context (one of utilization) to refer to the "CMA filtered waveform" as the final product of these steps. More strictly speaking, the "CMA filter" would refer to the algorithm employed in the first step of this process, determination of the baseline, and the thusly determined *baseline* would be the CMA filtered waveform: we will instead make reference to this step and its product only with the "baseline" context made clear.

The response of the CMA filter is demonstrated for two real signals from the scattering crystal for the CsI[Na] experiment in Figure C.2.

C.4.1: Description

Described more concretely, the CMA filtering with a filter length L of a waveform of length N proceeds as follows.

Preprocessing A rejection threshold and approximate baseline value are determined. Here, this step is executed using the "normal mode approximation" described in Sec. C.4.2.

Preload The first $L/2$ samples of the filter kernel are populated with the preload value determined in preprocessing.

Initial population Loop over the first $L/2$ samples of the waveform. For each sample w_i, compare with current average value of filter kernel μ_0. If $|\mu_0 - w_i|$ is less than the rejection threshold, add w_i to the sample kernel and update μ_0 based on new kernel contents.

Main loop over waveform For each index i in $i = 1, \ldots, N - 1 - L/2$, examine the new waveform sample at the leading edge of the filter range, $w' \equiv w_{i+L/2}$. If $|\mu_{i-1} - w'|$ is less than the rejection threshold, add it to the head of the filter kernel.

Update kernel and mean Check that kernel length is $\leq L + 1$; if larger, remove tail value. Determine μ_i, average of values in filter kernel.

Loop over final samples For the final $L/2$ samples of the waveform, remove the tail value of the filter kernel and then determine μ_i as the average of the remaining kernel.

For an implementation of this algorithm in `C++`, relying on `ROOT` [61] and `TWaveform`[3] [178], see Section E.1.

The filter behavior is potentially slightly different over the first and last $L/2$ samples of the waveform. In the former case, the preload value should ensure stability similar to that over much of the waveform, while in the latter the decreasing length of the filter kernel could make the filtered waveform slightly more subject to volatility of the baseline.

C.4.2: Determination of approximate baseline and noise values

Two parameters of the CMA filter require particular attention, and their calculation is discussed here. These parameters are:

Preload value The preload value is intended to stabilize the filter response in the presence of signal early in the waveform; to serve this purpose effectively, the preload value should be close to the average baseline value[4] of the first $L/2$ samples of the waveform.

Reject threshold Stepping to index i, the algorithm checks the sample at index $i + L/2$; if the difference between this new sample and the filtered value at index $i - 1$ exceeds the rejection threshold, the new

[3]The version of `TWaveform` actually used is based on a fork of the publicly available GitHub project. The forked version, `TWaveform-casa`, is intended for the COHERENT Collaboration and has some non-substantive modifications by G.C. Rich.

[4]Given time dependent behavior, "the average baseline value" over the length of the entire waveform *may* be different than that of the first few samples; determination of the average baseline over the entire waveform can also be somewhat more time consuming for waveforms with many samples.

sample is *not* added to the filter. The rejection threshold should be chosen such that it preferentially rejects samples including signal while accepting samples which represent noise or baseline fluctuation.

For the analysis of the QF measurements here, the baseline and noise fluctuations were assumed to be represented modestly well by a normal distribution. We then seek to determine the parameters μ and σ defining this shape.

It is computationally demanding to perform an actual fit to baseline data on an event-by-event basis. Additionally, as these parameters only serve to inform subsequent filtering, highly-precise results are not necessary. To establish approximate values for the Gaussian baseline shape with relative speed, the following algorithm is used:

Histogram X samples from waveform Some number X samples populate a histogram. X should ensure statistical fluctuations are not significant in the resulting histogram, but does not need to equal the length of the waveform.

Determine maximum bin The bin with the highest content is determined; this corresponds to the mode of the samples, which is a coarse estimate of the baseline and less sensitive to large deviations (i.e., signals) than a mean of the samples.

Bin weighted average to find μ Perform a bin-weighted average around the mode, which determines the approximate baseline mean μ in our Gaussian approximation.

Locate $\pm 50\%$ content bins Identify the last bin whose content is $\leq 50\%$ and the first bin whose content is again $\leq 50\%$ of that of the maximum bin.

Determine FWHM and σ The distance between bin centers of the $\pm 50\%$ bins determined above provides an approximation of the full width at half maximum (FWHM) of our Gaussian baseline approximation. For a more typical parametrization in terms of σ, use the relationship FWHM $\approx 2.355\sigma$.

With parameters for the Gaussian approximation of the baseline of the start of the waveform, we can provide informed inputs to the CMA filter. The mean μ of the Gaussian is supplied as the preload value, while the rejection threshold is specified to be $4\sigma \approx 4 * \text{FWHM}/2.355$.

Section C.5: Toy Monte Carlo waveform generation

The purpose of the toy Monte Carlo (TMC) event generator is to inform an accurate understanding of the effects of any signal processing techniques applied to the collected experimental data. Such an understanding is best established by the use of understood, controlled test signals which faithfully represent

the characteristics of real signals in all ways, including energy, timing, and any correlations which might exist.

C.5.1: Single photoelectron shapes

The shape of SPEs can be approximated well by the convolution of a Gaussian with an exponential decay [5]. Such a shape has a closed analytical PDF described by

$$f\left(t;\sigma,\tau,t_0\right) = \exp\left(\frac{\sigma^2}{2\tau^2} - \frac{t-t_0}{\tau}\right)\left[1 - \operatorname{erf}\left(\frac{\sigma^2 - \tau\left(t-t_0\right)}{\sqrt{2}\sigma\tau}\right)\right], \tag{C.1}$$

where σ is a standard parameter of the constituent Gaussian, t_0 is the mean of the Gaussian, and τ is the decay time of the exponential convolved with the Gaussian shape.

Data sheets and product manuals from Hamamatsu, the manufacturer of the PMT utilized in these experiments, provide some information on timing characteristics that might be expected of SPEs. The H11934-200 ultra-bialkali PMT used to observe the CsI[Na] crystal has approximate rise and fall times of 1.3 ns and 5.8 ns respectively when operated at -900 V bias [199, 208], slightly lower than the -950 V bias used for these measurements. Each of the photoelectrons identified via thresholding in the pretrace of waveforms from a single hour-long run were fit with a function consisting of the shape defined in Eq.(C.1) and a simple constant (baseline) offset. The fits were carried out only over a localized region of the waveform, including only 100 ns before and after the threshold crossing.

Only a small subset of collected data from the CsI[Na] measurement, collected over the course of a single hour, was used to inform representative shape parameters, The timing parameters used in the toy MC generation are fixed, rounded values informed by the fit to the parameters as described above; therefore, fluctuations in the timing characteristics of the SPE shapes are not reflected in the toy MC. Beyond fluctuations, it is not expected that the timing characteristics of single photoelectron pulses will appreciably change over the timescale of several days in the absence of extreme circumstances. Use of a limited range of data may also underestimate the impacts of PMT gain drift and variations in the SPE amplitude. With the simple parametrized model ultimately used to characterize the filter response in Section C.6.3, small variations in timing and SPE gain should be negligible.

C.5.2: Distribution of photoelectrons

For every event generated by the toy MC routine, the number of photoelectrons present is an unrealistic abstraction from physically realizable configurations where systematic and statistical fluctuations are largely

[5] In `ROOT` and specifically `RooFit`, this shape is easily accessible as a single-sided `RooDecay`.

unavoidable. Each generated event contains a strictly specified number of photoelectrons. This number is *not* subjected to any statistical fluctuations: each time an event with one PE is requested, a generated waveform with a single PE is returned. Consequently, care must be taken if the toy MC routine is to be used to model very specific experimental circumstances, where energy-dependent scintillation yields, light collection effects, and statistical spreading must all be taken into account when specifying the number of PEs to include in the model waveforms.

The timing distribution of photoelectrons in the generated waveforms accounts for variation in the light production time of the CsI[Na] scintillator as well as the (comparatively small) photoelectron transit time in the PMT. Compared to the characteristic time of a single photoelectron signal from the H11934-200 PMT, the scintillation photons produced in the CsI[Na] crystal are distributed over a relatively long time. Scintillation light production has been measured by several earlier efforts, and though the fitted parameters differ, the distribution is typically modeled by an additive combination of two exponential decays, with fast and slow time constants on the order of a few hundred nanoseconds and a few microseconds. Parameters used here come from Collar *et al.*, who report a fast decay time of $\tau_{\text{fast}} = 589 \pm 4$ ns, a slow decay time of $\tau_{\text{slow}} = 6.7 \pm 2.4$ μs, and a ratio of the signal intensity between the two of $I_{\text{slow}}/I_{\text{fast}} \sim 0.41$ [77].

C.5.3: Baseline and noise "farming"

The principal motivation for a careful choice of filtering approach is the removal of subtle variations in the baseline which do not integrate to 0 through use of a mode-type baseline determination algorithm. If crucial properties of the baseline and noise were well defined and known *a priori* then a targeted filter could be employed such as a high-pass filter intended to remove slowly-varying, 60-Hz line noise or a band-stop filter addressing noise of a specific frequency. In the case of the CsI[Na] experiment, however, the spectral distribution of the baseline and noise are not confined to one, or a few, narrow band(s), and the higher-frequency components overlap with the signature of SPE-like signals.

The legitimacy of any investigation of filtering that relies on artificial, generated event waveforms is predicated on accurate representation of the experimental baselines. Rather than attempting to derive a parametrized approximation of the baseline and noise structure, the data itself is called upon to provide empirical input. Since the trigger of the DAQ does not place any requirement on the presence of signal in the scattering detector (see Section 4.2.2), there are digitized events which include no photoelectrons in the scatterer signal.

These prototype events are identified by running the waveform through a level threshold algorithm, where the threshold is defined relative to a baseline determined by a mode-type approach. This technique could fail

to identify hypothetical low-amplitude signals, including them in what is considered a sample of waveforms without any signal.

When the generator routine is invoked, a "reference" data file from which baseline samples will be drawn is specified[6]. For each event to be generated, the generator routine samples a random number between 0 and the number of events contained in the data file supplied as a reference; the event at the randomly chosen index is checked for the presence of photoelectrons; if the randomly chosen event is found to have any non-zero number of PEs, another random number is chosen; the process repeats until an event with 0 PEs is found and this 0PE-waveform is retrieved from the reference file for use as the template upon which the simulated event is assembled. At the entrance of the event-generation routine a random number seed is determined based on the system timer, ensuring[7] that subsequent event generation will not immediately select the same baseline waveform.

Section C.6: Evaluation of CMA filter using toy MC waveforms

C.6.1: Goals of the tests

As the filter will directly impact the amplitude of the signals extracted from the experimental data, care must be exercised so any possible bias introduced by the filter is understood and accounted for in analysis. For reference, Fig. C.3 shows SPE "charge" distributions from the toy MC before and after the CMA filter is applied; the toy MC allows a comparison between the "original" signal and that which has been acted on by the filtering algorithm. The primary goal of these tests is to establish an understanding of the relationship between the integrated signal from a filtered waveform and an idealized, controlled input signal across a range of test signal "energies".

C.6.2: Comparisons to be made

Prior to any potential filtering effects, the digitized signal is an imperfect representation of the "actual" signal. A direct comparison between the "actual" or "analytical" integral of the input signal(s) and the filtered integrals does not unambiguously expose the contribution of the filter to any discrepancies between the two values as such a comparison implicitly includes the additional impact of discrete-time sampling of

[6] Along with the reference file, a file containing a preprocessed version of the same data is supplied. The preprocessed file contains, among other information, the number of PEs counted by the level threshold algorithm for each waveform. Use of the preprocessed file to quickly determine if an event has no PE content substantially reduces the time required to find a suitable baseline sample.

[7] More properly, the timer-based seed creation *considerably decreases* the probability that subsequent requests for baseline waveforms will proceed through the same sequence of reference events.

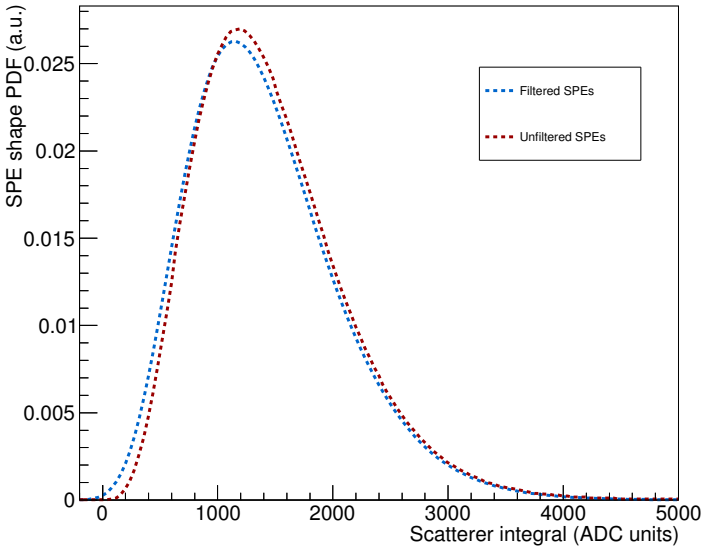

Figure C.3: SPE "charge" distributions from the toy MC model with and without the CMA filter applied. Minor differences are apparent. The intent of the toy MC analysis is to determine corrections which can be applied, on a statistical basis, to the measured integrals of filtered waveforms.

continuous signals. In certain applications these digitization effects could be of significant interest, and while they are not under direct investigation here, they must not be inadvertently conflated with filtering effects.

A comparison between the "discretized" integral calculated using the analytical PE model and the filtered integral provides the most direct reflection of the effect of the filter on the extracted signals.

C.6.3: Assessment of CMA filter response

Using the same waveform processing algorithms (including, for instance, the same integration routines) as are utilized in the primary data analysis, waveforms generated by the toy MC code with and without "farmed" noise were integrated and the resulting values compared. Events were generated with a wide range of photoelectron content, ranging from 1 PE to 1000 PEs. For each number of PEs, 1000 unique events were created by the toy MC.

The distribution of the ratio between the integrals of the noisy and noiseless waveforms was approximated with a gaussian fit. This approximation is imperfect, especially at lower PE numbers where the distribution is more sharply peaked; as the gaussian tends to overestimate the width of the distribution where deviation from normality is apparent, this will provide a more conservative estimate of the performance of the filter in faithful reproduction of the signal integrals. Fitting of the ratios resulted in parameter estimates for both the mean μ and width σ of the response model, along with corresponding error estimates for each parameter.

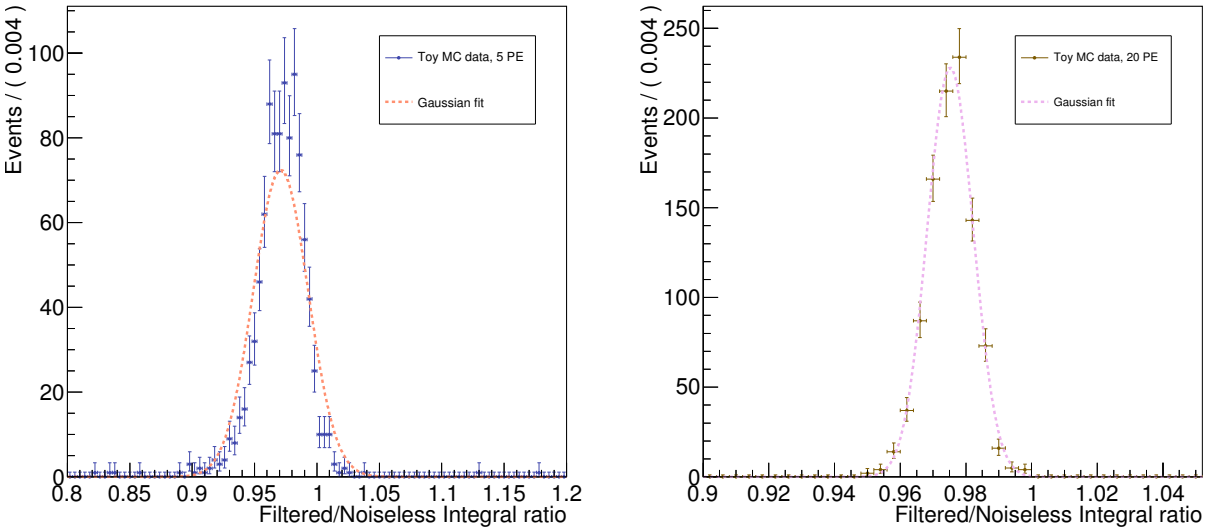

(a) Ratio of filtered to noiseless integrals for 5-PE events. (b) Ratio of filtered to noiseless integrals for 20-PE events.

Figure C.4: Distributions of the ratios of integrals from filtered, noisy waveforms and unfiltered, noiseless waveforms for 5-PE and 20-PE events. The waveforms were generated by the toy Monte Carlo described in Section C.5. Also shown are fits to these distributions assuming a gaussian shape, providing only a coarse approximation of the shape for events with low PE numbers where the gaussian model is less sharply-peaked than the data, tending to overestimate the width of the distribution.

Example ratio distributions and fits can be seen in Figure C.4. The parameters describing the gaussian approximation of the impact of the filter on waveform integrals over a range of photoelectron content can be seen in Figure C.5 along with overlaid fitted models of the PE-dependent trend.

As both the mean and width of the response shape are of interest, simple parametric models for each quantity as a function of photoelectron content were developed. The parameters were fit in a single Markov chain Monte Carlo (MCMC) routine, but the models were independent and compared solely to the value and error data corresponding to the parameter with which the model pertained. Models for the response parameters were parametrized as

$$\mu\left(n_{\mathrm{PE}}\right) = \mu_0 + \mu_m \log_{10}\left(n_{\mathrm{PE}}\right), \tag{C.2a}$$

$$\sigma\left(n_{\mathrm{PE}}\right) = \sigma_0 \left(n_{\mathrm{PE}}\right)^{\sigma_a} + \sigma_\Delta, \tag{C.2b}$$

with floating parameters μ_0, μ_m, σ_0, σ_a, and σ_Δ. These models were fit to the data shown in Figure C.5. Though the models shared no explicit interdependence, a joint likelihood between the models for μ and σ was constructed and used to guide the MCMC. A corner plot showing the distributions of the parameters determined by the MCMC fit and their covariances can be seen in Figure C.6; parameter values are also shown

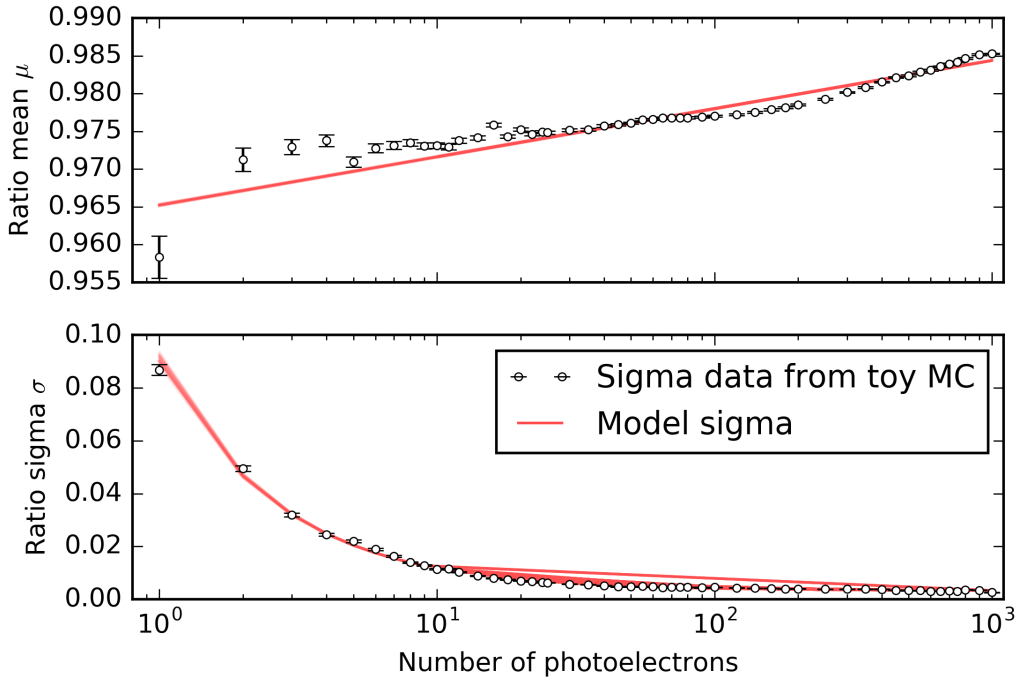

Figure C.5: The values of the mean μ and sigma σ parameters describing a gaussian model of the CMA filter response as a function of the number of photoelectrons present in the waveform. The filter response is defined here as the ratio of the integrals from a noisy, filtered waveform and a noiseless, unfiltered waveform; examples of these ratio distributions and the corresponding fits can be seen in Figure C.4. Each data point, corresponding to a prescribed number of photoelectrons, included 1000 waveforms generated by the Monte Carlo model described in Section C.5. Section C.6.3 of the text discusses the model in greater detail.

Parameter	Fit value
μ_0	$9.65E-01^{+9.29E-05}_{-9.05E-05}$
μ_m	$6.39E-03^{+3.75E-05}_{-3.91E-05}$
σ_0	$8.75E-02^{+1.32E-03}_{-1.32E-03}$
σ_a	$-1.01E+00^{+7.84E-03}_{-7.37E-03}$
σ_Δ	$3.23E-03^{+1.94E-05}_{-2.04E-05}$

Table C.1: Parameters describing the PE-dependent response of the CMA filter, characterized as the ratio of the integrals of the filtered waveform and the unfiltered, noiseless waveform. The model is described by Equations (C.2). The covariances of these parameters are shown graphically in Figure C.6.

in Table C.1. Figure C.7 shows the photoelectron-dependent impact of the CMA filter on the determined integral along with a sample of the mean and $\pm 1\sigma$ regions from the MCMC fit, shown in red.

Section C.7: An alternative approach based on interpolation

The filter discussed here is similar in spirit to that of Abbiati *et al.* [2] which shares the goal of providing a measure of a non-stationary baseline. Abbiati and coauthors pursue an alternative approach based on interpolation between baseline samples prior to and following a pulse; they conduct a rigorous exploration of their filter's response in the frequency domain which is not undertaken here for the CMA filter [2]. They specifically mention the kernel of the approach taken here, "stop the mobile average", as the other of two options, but they point out that this path precludes the use of mathematically optimal filters as the estimation will be potentially corruptible by noise [2].

While a thorough response to the assertion made by the authors that a moving-average based filter is corruptible by noise is not offered here, it should be noted that the CMA algorithm is not assumed to be mathematically optimal. The numerical performance assessment presented in Sec. C.6.3 using the toy MC tools of Sec. C.5 is sufficient to provide confidence in the filter and its application here.

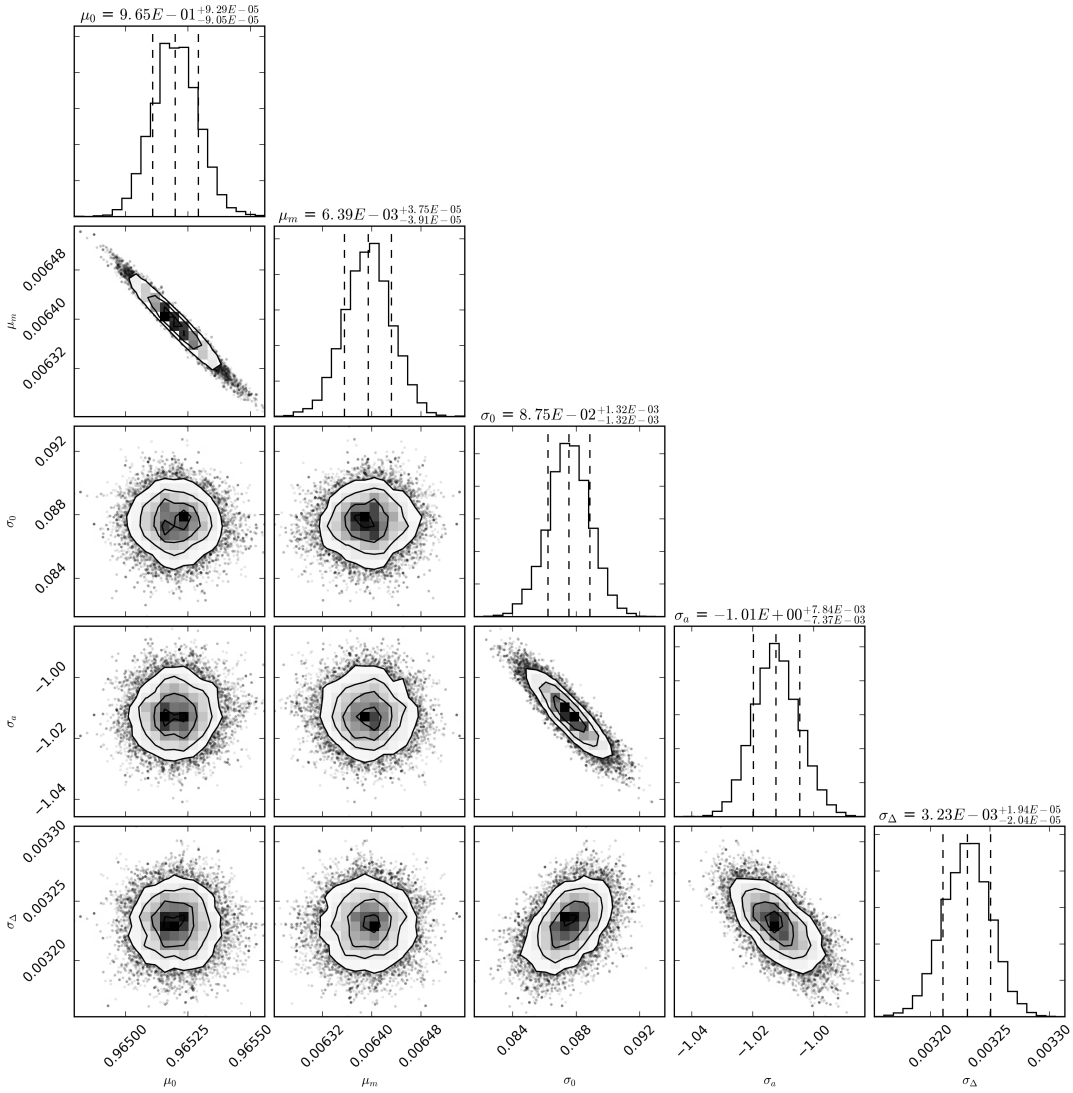

Figure C.6: Parameter distributions and covariances associated with a model of filter response as a function of the number of photoelectrons present in a waveform. These distributions are the product of an MCMC fit with a main chain of length 1500 steps following a burn-in of 10k steps, though the parameters largely converge on their final values within the first ∼500 steps. Section C.6.3 of the text discusses the model in greater detail.

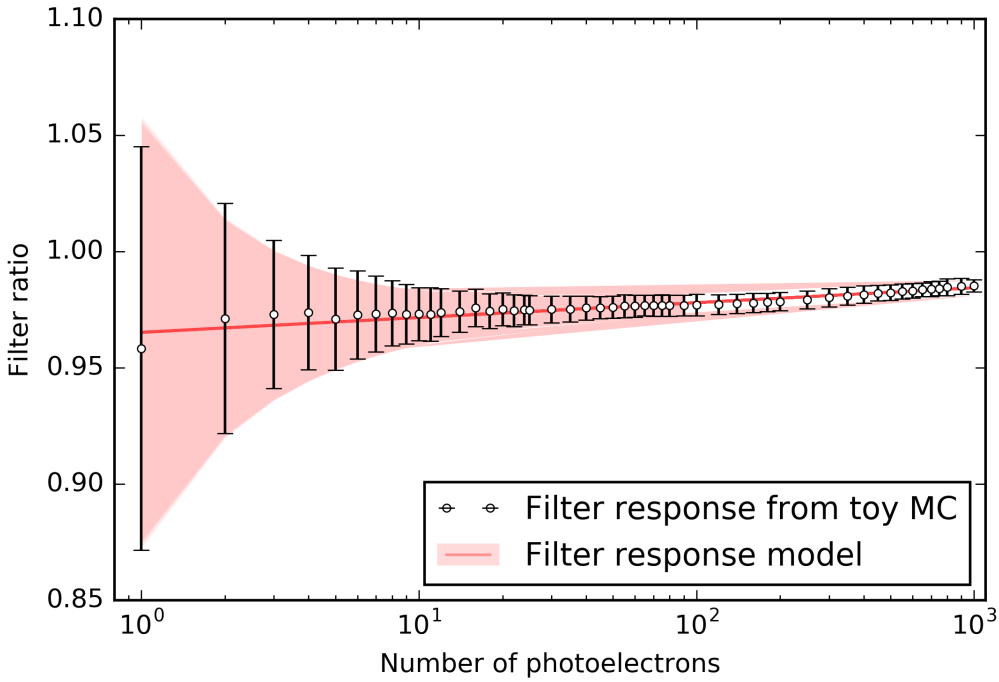

Figure C.7: The ratio of the integral of a noisy waveform subjected to CMA filtering and a noiseless, unfiltered reference waveform. Data points are from 1000 unique events per PE content; these events are generated by the toy MC process described in Section C.5. The distribution of the ratio of these integrals is modeled as a gaussian, and the values of μ and σ of these gaussians define the y-values and the error bars of the data points, respectively. The PE-dependent response is modeled as described in Section C.6.3 and the red shaded region shows $\pm 1\sigma$ from the resulting model, with the more intense red line showing a sample of the mean values from this model.

APPENDIX D: MARKOV CHAIN MONTE CARLO TIME-OF-FLIGHT FITTING

Section D.1: TOF data collection and preparation

Time-of-flight data is collected by the "zero degree" monitor detector as described in detail in Sec. 4.3.1. The collected data represents the observed energy deposition in the zero-degree detector, the corresponding pulse-shape parameter, and the timing relative to the beam-pickoff monitor all correlated on an event-by-event basis. Raw data from the CODA DAQ software [133] is stored in a custom format but for processing the data is converted[1] from the CODA-style output file into a ROOT TTree [61].

As a final step in preparation for the fitting, the TOF spectra are baseline subtracted and TOF regions-of-interest (ROIs) are identified for each standoff distance of the zero-degree detector. These ROIs are selected following baseline subtraction such that they contain a large fraction of the primary neutron population; it is also assumed that the primary neutrons are the only source of counts in the spectra.

Section D.2: Neutron energy from time of flight

Extraction of neutron energy from time of flight relies on the determination of the time at which the neutrons were produced t_0 and time at which they were detected t_D, along with a measurement of the distance the neutrons traveled during this interval. This knowledge allows reconstruction of the neutron energy from simple kinematic relations[2]. In an actual analysis, this idealized treatment must be reconciled with experimental realities, including finite and non-zero widths of timing and energy distribution in addition to extended experimental geometries. Further complication of the recorded timing spectra is added by the effects of charged-particle transit time and energy loss across the neutron production volume[3].

For the measurements of Chap. 4, neutrons were generated in an extended production target taking the form of a ∼1-1/8"-long gas cell discussed in more detail in Sec. 4.1.2. In this configuration, the incident deuteron beam loses a non-negligible amount of energy along the length of the cell resulting in neutrons of different energies being produced at different locations, further complicated by changes in the $\mathrm{D(D},n)^3\mathrm{He}$ cross section over the range of relevant deuteron energies. Additionally, transit of the deuteron through the cell takes a finite amount of time which must be considered along with the differing flight-path lengths

[1] The conversion software was developed by Ron Malone.

[2] At the energies involved in these experiments, relativistic effects can be neglected.

[3] In the case of some quenching factor measurements conducted during the course of this research, but not discussed in this thesis, the $^7\mathrm{Li}(p,n)$ reaction was used produce neutrons in a thin LiF neutron-production target with thickness ∼500 nm: this somewhat simplifies the analysis as the incident proton beam loses little energy through the entire production volume which can be considered infinitely thin.

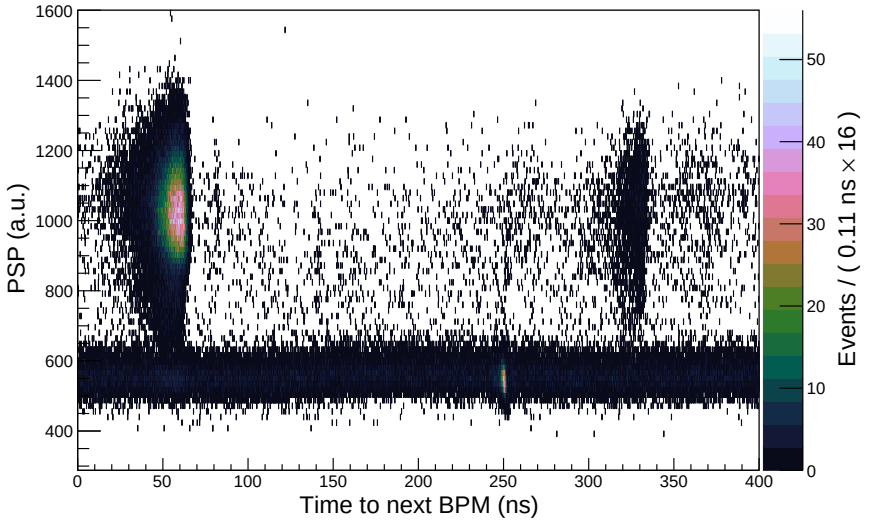

(a) Two-dimensional view of TOF data

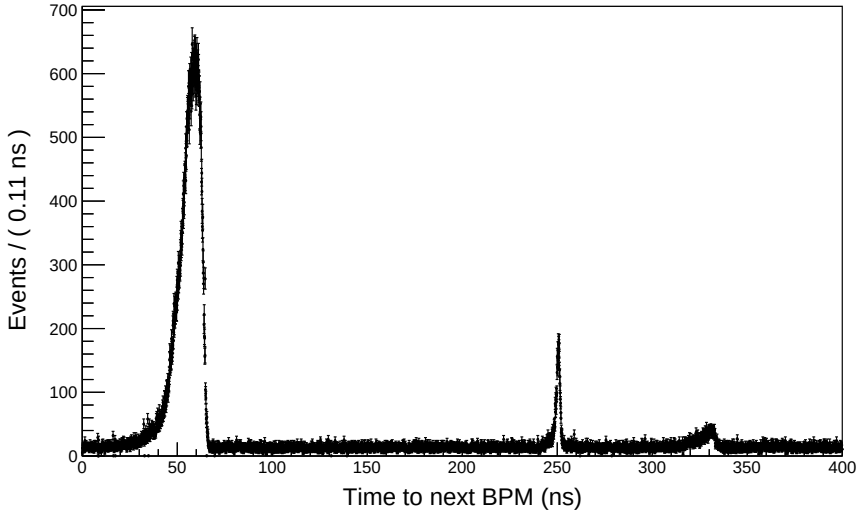

(b) Projection onto TOF-only axis

Figure D.1: Time-of-flight and PSD data from the zero-degree beam monitor. Figure (a) shows the PSP as a function of separation in time between the detector signal and the next BPM pulse, while (b) shows this data projected onto the time axis. A very prominent, high-PSP feature at ∼50 ns is the primary neutron population, while the low-PSP population at ∼250 ns corresponds to the γ-rays produced by the charged-particle beam as it strikes material in the beam line. The region in time around 250 ns is the "γ-timing" region fit with a model that serves as a TOF calibration; see the fit to this region in Figure D.2.

of the neutrons produced at different points along the cell and the relative velocities of the two particles. Correlations exist between all of these factors and must be taken into account.

Section D.3: Analysis of γ-ray features and determination of t_0

Since time-of-flight energy analysis is based on the time $\Delta t \equiv t_1 - t_0$ over which a particle travels a distance ℓ, we must produce measurements of both the start t_0 and stop t_1 times of an interval associated with a well-understood neutron flight path. As the charged particle beam at the TUNL tandem accelerator can be pulsed, a natural choice is to use the timing of the incident, charged-particle beam pulse as t_0 and the time of interactions in a detector positioned directly in the produced beam as t_1. Making use of the fact that the charged-particle beam produces γ-rays in addition to neutrons, the locations of the γ-ray peaks in collected TOF spectra can provide a convenient mechanism for determination of t_0 and provide insight into the timing distribution of the charged-particle beam.

The beam of charged particles from the tandem accelerator can produce γ-rays through interaction with any component along the beam line, not only those in the neutron-production volume[4]. To provide a reckoning of t_0, a γ-ray population associated with a confidently known production site must be selected: the distances between the γ-ray production site, the neutron production site, and the detector will factor into the TOF analysis. Examination of the collected TOF data for the CsI[Na] measurements shows three recognizable γ-ray populations in a relatively small region of TOF space; the associated production sites are assumed to be (in order along the direction of the beam): the collimator before the gas cell, the havar foil on the entrance to the gas cell, and the tantalum disk at the end of the gas cell.

The γ-rays will arrive in this order on the detector and the observed intensities of the populations suggest that the havar foil is the most prominent source of γs: this is consistent with expectations, as the deuterons incident upon the foil should have lost little energy prior to this interaction and the beam is tuned to maximize flux on the gas cell target whose entrance is subtended by the foil, thereby minimizing the current on the collimator. No significant attenuation of the beam flux is expected through the havar or the deuterium gas, so a large fraction of the deuterons incident upon the havar will also reach the tantalum disk at the end of the gas cell; though the flux on this disk may be nearly equivalent to that on the havar, the considerable deuteron energy loss expected through the foil (and the additional, modest loss through the gas volume) should result in a significantly lower yield of inelastic γ-rays from interaction with the W disk relative to the havar.

[4]Generally speaking, neutron production is not necessarily restricted to a single volume, either, but the relatively-low proton and deuteron energies used for these experiments significantly limit the locations of neutron production by virtue of being below threshold for many neutron-production reactions aside from those in the intended production volume.

The region of the TOF spectra with the 3 near-target γ-ray populations is analyzed using the RooFit fitting package [247] in ROOT. The model used for the spectra consists of 3 peak shapes additively combined with a uniform distribution modeling a flat, accidental background; the peak shape used[5] is that of a Gaussian with mean μ and width σ convolved with an exponential decay with time constant τ. Using the assumption that each of the γ-ray populations are produced similarly (specifically, that they all are produced in an infinitely thin target and the deuteron transit time over this region is negligibly small), the parameters corresponding to the relative timing features of these peak shapes are common between the three. More explicitly, if the distribution in time for peak i, in $i = 1, 2, 3$, is defined by parameters μ_i, σ_i, and τ_i, we fix

$$\sigma_\gamma = \sigma_1 = \sigma_2 = \sigma_3, \qquad \tau_\gamma = \tau_1 = \tau_2 = \tau_3,$$

only allowing the absolute timing of each pulse μ_i to vary independently. An extended maximum likelihood (EML) fit is carried out using this model over the γ-ray TOF region, and the means of the Gaussians involved in the peak shapes, μ_i, are taken to represent the time $t_{\gamma,i}$ of the γ-ray time of flight. In this fit, count-rate normalization parameters (i.e., the "extended" component of the EML fit) are allowed to float without constraint and are ultimately immaterial to the result, aside from qualitative interpretation of the relative magnitudes. The shape characteristics of each peak, whose distribution is governed by the parameters σ_γ and τ_γ, are interpreted as a representation of the timing distribution of the deuteron beam itself.

The most intense γ-ray peak, corresponding to production at the havar foil, is selected as a reference; with the location of production thusly identified, the flight path ℓ_γ from production site to detection site can be determined. Writing the mean of the reference peak as μ_{ref} and neglecting the transit time of γ-rays across the monitor detector, we can express the *charged-particle-beam arrival time*,

$$s_0 = s_{\text{ref}} - \frac{\ell_\gamma}{c}\phi_{\text{cal}}, \qquad s_{\text{ref}} = \mu_{\text{ref}}, \tag{D.1}$$

where we must be careful to note that we are working in digitized unit space, utilizing the calibration ϕ_{cal} with units of channel / nanosecond.

The timing distribution of the charged-particle beam is taken into account when fitting neutron timing data.

[5]This shape, a Gaussian convolved with an exponential decay, is a `RooDecay` in the parlance of RooFit.

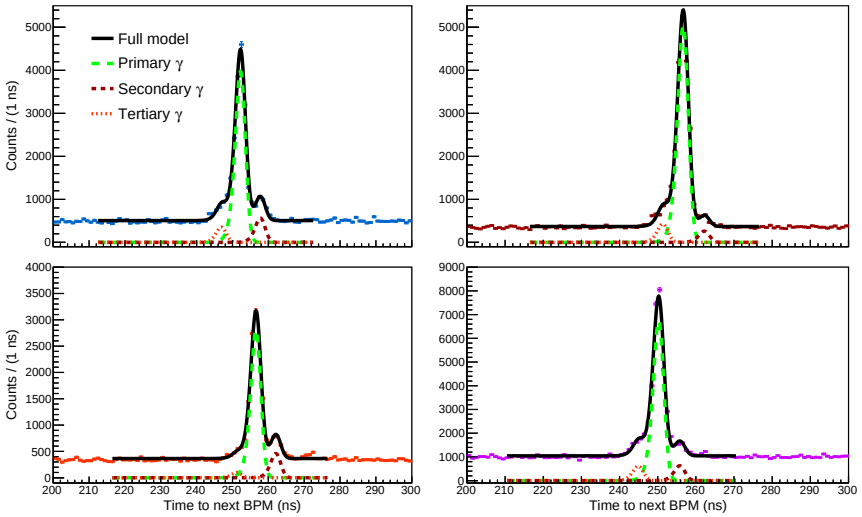

Figure D.2: Clockwise from upper left: middle standoff; close; far; close detuned. The black line in each plot shows the total timing PDF fitted to the data; the green, coarse-dashed line shows the "primary" γ-ray population associated with the Havar foil; and the red-orange, dotted and crimson, fine-dashed lines show populations associated with the tungsten disk at the end of the gas cell and the collimator before the gas cell, respectively. The relative prominence of the collimator population in the close detuned run compared to the close run provided confidence in the identification of the source of beam-correlated γ-rays along the beam line by intentionally directing more beam current onto near-target apertures (see discussion in Section D.3)

Section D.4: Physical model for TOF spectra

To develop an accurate model of the observed TOF spectra, consideration must be given to the physical processes involved in neutron production. Ultimately, the distribution that the fitting routine produces, which we call $f_0(E_D)$, corresponds to the *energy distribution of deuterons after they have passed through the havar foil, as they enter the deuterium gas volume.* This shape is modeled by a shifted, additively-inverted log-normal distribution,

$$f_0(E_D) \equiv E_{\text{beam}} - f_{\text{LN}}\left(E_D; \vec{\theta}\right), \tag{D.2}$$

where the floating parameters are E_{beam} and the entries in the parameter array $\vec{\theta}$ which parameterize the log-normal distribution f_{LN}. The conceptual basis for this shape is that the deuteron beam incident on the havar foil is monoenergetic and has some energy loss in the havar foil: the E_{beam} parameter corresponds to the energy of the incident beam and the energy loss through the havar is modeled by the log-normal distribution.

Energy loss through the havar foil would be more rigorously modeled by a Vavilov distribution [169] which introduces considerable computational complexity. Approximate treatment of the energy-loss distribution is realized as a log-normal distribution, described by a shape parameter σ, a location parameter θ, and a scale

parameter m, taking the form [189]

$$f_{\mathrm{LN}}(x;\theta,\sigma,m) = \frac{\exp\left(\dfrac{-\left(\ln\left[(x-\theta)/m\right]\right)^2}{2\sigma^2}\right)}{(x-\theta)\,\sigma\sqrt{2\pi}}, \qquad x>\theta; m,\sigma>0. \tag{D.3}$$

Specifically, the TOF model developed here accounts for:

Deuteron transit A deuteron incident on the production cell takes a non-zero amount of time to transit across the cell to the site where a neutron is produced. This is approximated by assuming all deuterons have an equal, representative, and non-changing kinetic energy E_{eff}; the time to arrive at the neutron production site, distance x_0' from the entrance to the cell, is then simply $t_0' = x_0'\sqrt{m_n/2E_{\mathrm{eff}}}$.

Deuteron straggling As the deuteron travels across the cell, it loses energy in the gas. The stopping power for the deuterium gas is calculated using the Bethe formula; the value used to represent the mean excitation energy for deuterium is 19.2 eV [125] and a density of 8.565×10^{-5} g/cm^3 is used in the calculation to represent the nominal pressure of 0.5 atm.

D(D, n)^{3}He cross section Deuteron-energy dependence of the production cross section was accounted for by using a cubic spline interpolation of the tabulated 0° cross section data points of Liskien and Paulsen [173].

D(D, n)^{3}He neutron energy As the deuterons lose energy, the energy of the produced neutrons is determined using Eq. (4.2).

Neutron flight time The time associated with flight of a neutron, with energy E_n, from the point of production to the face of the beam-monitor detector (for convenient short hand, a distance x_n) is calculated by $t_n = x_n\sqrt{m_n/2E_n}$.

Neutron detection time Spread in detection time (and observed TOF) associated with neutron transit across the detector is approximated by convolution with an exponential.

There are several approximations or assumptions involved in the use of this model. The experimental geometry is treated solely in the single dimension of the beam axis and all units of length are considered to exist entirely in this single dimension; no isotropic effects are considered. Similarly, neutrons are assumed to be produced only along the axis of the incident deuteron beam which is assumed to be colinear with the the neutron beam axis. For small angles around 0°, there is little change in either the produced-neutron energy or the differential cross section for the D(D, n)^{3}He reaction [173]; since the geometry of the SSA aggressively selects very-forward-produced neutrons, this simplified treatment of production is well justified.

Section D.5: The Markov chain Monte Carlo method

To briefly introduce the concept of Markov chains and the technique of MCMC fitting, we consider a scenario where we have a model with a single parameter θ. Succinctly, a Markov chain in θ is a series of numbers $\theta_0, \theta_1, \ldots, \theta_n$, where the value θ_{n+1} depends *only* on the value of θ_n [188]. In a Markov chain, the transition probability $T(y, y')$ describes the probability of θ_{n+1} assuming a certain value y' given that the value of θ_n is y. There are a few important properties of chains to note: if the transition probabilities are independent of time[6], a chain is homogenous; a value θ_Ω for our parameter is said to be stationary if, once it is reached, it persists; and a chain is ergodic if $\theta \to \theta_\Omega$ irrespective of the initial value of θ as the number of steps in the chain tends towards infinity [188].

In performing MCMC fits, the task is to sample from many Markov chains of our parameter θ. For each step in the chain, the transition probability is informed by our model evaluated with parameter θ and its compatibility with the data that is being fit. By stochastically sampling from many chains over many steps, the values of θ contained in the chains should represent a stationary distribution of the value of θ which best describes the data, given the chosen model.

The process of determining the samples in a Markov chain is a subject with a long history beginning with the work of Metropolis *et al.* in 1953 [184]. The Metropolis [184] and Gibbs [117] sampling algorithms are canonical examples about which many introductions have been written. While exposition on these processes is left to the literature, the approach toward drawing a new parameter value θ_{n+1} is generally stochastic[7] with some pressure applied towards values that better reproduce the observed data using the given model.

The affine invariant sampling scheme introduced by Goodman and Weare [123], building off of work by Christen and Fox [70], is a powerful approach that can very efficiently explore a parameter space and dramatically speed up convergence. Rather than using independent chains, affine invariant sampling utilizes an ensemble of "walkers" which, during the step proposal process, use information on the location of other walkers to adaptively improve the sampling of the parameter distribution. This adaptive approach allows efficient sampling of pathological distributions and does not require *a priori* knowledge or specification of parameter correlations [123]. An implementation of affine invariant sampling is a prominent feature of the **emcee** Python package, which provides a framework and set of utilities upon which MCMC calculations can be built [107, 108].

[6] To be perhaps needlessly clear: this means that $T(y, y')$ will be the same when evaluated tomorrow as it is when evaluated today; it is not related to whether or not the model has some dependence on a "time" variable.

[7] The stochastic nature of the sampling process is responsible for the "Monte Carlo" component of the name for this technique.

Section D.6: Implementation and running of MCMC TOF fitting routine

The MCMC fitting process requires the ability to evaluate a probability distribution and compare it with available data at each step in parameter space. Unfortunately, there is not a closed-form expression for the TOF that results from the physical model described in Sec. D.4 is not readily established. Unable to analytically evaluate the probability at each TOF value, a high-statistics toy dataset is drawn from the model for each set of parameter values and a likelihood is calculated by comparing an appropriately renormalized, binned version of this data against the experimental TOF distribution.

Data was generated from the model in the following steps:

1. Draw n samples from the initial energy distribution $f_0(E_D)$. Let E_μ be the mean of the drawn samples.

2. For each sample, use the stopping model to calculate the energy at each bin center along the length of the gas cell. Increment bin contents of a 2-D histogram for each (E_D, x) pair, summing all of the n samples together. Normalize this histogram once all samples have been added.

3. For each bin center (E_i, x_j) in the normalized histogram, determine an effective deuteron energy E_D^{eff} by averaging E_i and E_μ. Approximate deuteron transit time t_D to this location, a distance x_j into the gas cell, $t_D = x_j\sqrt{m_D/2E_D^{\text{eff}}}$. Determine E_n^0, the energy of a neutron emitted at 0° from $\text{D}(\text{D},n)^3\text{He}$ by a deuteron with energy E_i. Approximate neutron flight time from this bin center to the face of the beam monitor, calling this distance x_n, $t_n = x_n\sqrt{m_n/2E_n^0}$.

4. Still for each bin center, sample a small time perturbation ϵ_t from from a neutron-energy dependent distribution associated with transit time across the beam-monitor detector prior to detection; drawn from a normalized PDF of the detection times, let w_t be the weight associated with the bin of the sampled ϵ_t value. Define a provisional TOF value $t^* \equiv \epsilon_t + t_n + t_D$.

5. Weighted by a factor of w_t, increment the bin associated with t^* of a 1-D histogram of TOFs from all bin centers.

6. After having processed all bins, convolve the 1-D histogram of provisional TOF values with a PDF of the deuteron-beam arrival time distribution, as found in Sec. D.3.

7. The histogram resulting from the convolution, scaled by a normalization parameter which is a floating parameter of the fit, represents the TOF distribution described by the model.

The TOF data generated by the above procedure is then used in the evaluation of a log-likelihood function conditioned on the experimental TOF data. For each bin i, an approximate continuous Poisson distribution

with mean $\mu = \mathrm{TOF}_i^{\mathrm{model}}$ is evaluated at $x = \mathrm{TOF}_i^{\mathrm{obs}}$. The logarithms of the resulting PDF evaluations for each bin are summed together to produce the log-likelihood value. The fit region is chosen to minimize the number of 0 or few-entry bins, but to avoid numerical problems, a bin with 0 contents in either the model or experimental TOF data is populated with a single count.

As discussed in Sec. 4.3.1, data was taken at several beam-monitor-detector standoff distances. Considered together, these measurements could provide data showing TOF shape differences that help disambiguate the neutron energy distribution underlying the observed TOF spectra. Slight modifications to the fitting process were made to include additional data: for each step in the Markov chain, model TOF data is generated for each of the 4 unique detector standoff distances and a simultaneous fit is conducted by utilizing a joint likelihood function (an additive combination of the log-likelihood values from model and experimental data for each of the standoff distances). The simultaneous fit to 4 sets of data based on collection of TOF spectra at 3 unique standoff distances from the gas cell was carried out using the Stampede supercomputer at the Texas Advanced Computing Center, part of XSEDE [241].

Section D.7: MCMC TOF fitting results

D.7.1: Model posterior

The output of an MCMC fitting procedure is a posterior distribution in the parameter space under investigation. In the experiments discussed here, this corresponds to a distribution in the parameters defining the *initial* energy distribution of the deuterons as they enter the deuterium gas cell $f_0(E_D)$ (see the discussion in Sec. D.4). Our model approximates transport of the beam through the cell under the assumption that the energy loss is sufficiently-well represented by Bethe stopping and that the cross section for neutron production as a function of deuteron energy is given by Ref. [173] with kinematics described by Eq. (4.2); more detail of the physical model is included in Sec. D.4. Using the same transport and production model as the MCMC fitter, the distribution $f_0(E_D)$ can be propagated forward to produce distributions of the neutron energy through the gas cell and of the expected TOF.

An MCMC "fit result" can be explored and visualized in several ways. Both the behavior of the model along the way and the result can be seen in a plot of "the chain", as in Fig. D.3, which shows the parameter values of the ensemble at steps along the chain. The burn in stage can be seen clearly in the early steps of Fig. D.3 where the walkers have been given initial locations normally distributed about a nominal "guess" for the parameter. This figure demonstrates one of the virtues of MCMC approaches: the ensemble explores the parameter space effectively, with excursions helping to ensure that, for instance, initial parameter values to not cripple the model with bias. Initial values for the parameters of this fit were chosen such that the

model evaluated cleanly (the walkers can explore parameters to the point where infinite results, or numerical errors, occur) and the TOF spectra yielded by the guesses were qualitatively close to the experimental values. The width of the gaussian defining the distribution of initial walker locations for the beam energy parameter (parameter 0) was $\sigma = 10$ keV as this parameter should be relatively well defined by experimental parameters (see discussions in Chapter 4.1). For other parameters, the width of the initial distributions were 5% (for the parameter corresponding, roughly, to mean energy loss through the havar foil), 10%, and 20%, with the final two parameters having less direct input from the experimental configuration. Figure D.4 shows a corner plot[8] of the parameters for the energy distribution.

Correlation of the parameters can also be explored visually with MCMC models, and this is especially convenient with affine-invariant samplers such as **emcee** where the correlations do not need to be known *a priori*.

D.7.2: TOF distribution from model posterior and comparison with data

Since the output of the fit is in terms of the deuteron distribution, comparison of its "results" with experimental data requires evaluation of the model. This comparison is referred to as a "posterior predictive check", and this can be seen for the TOF data in Fig. D.5, where the model with parameters fitted by the MCMC routine has been used to produce the line overlaid the collected TOF. Though disagreement with some of the TOF data, especially that at shorter standoff lengths, is apparent, the deviations correspond to energy differences which are small compared to the width of the overall neutron-energy distribution. The neutron energy distribution produced by the MCMC fit can be seen in Fig. D.6, where it is shown alongside the deuteron energy distribution determined at both the beginning and the end of the gas cell.

[8]This figure was made using `corner`, a Python package described in Ref. [106].

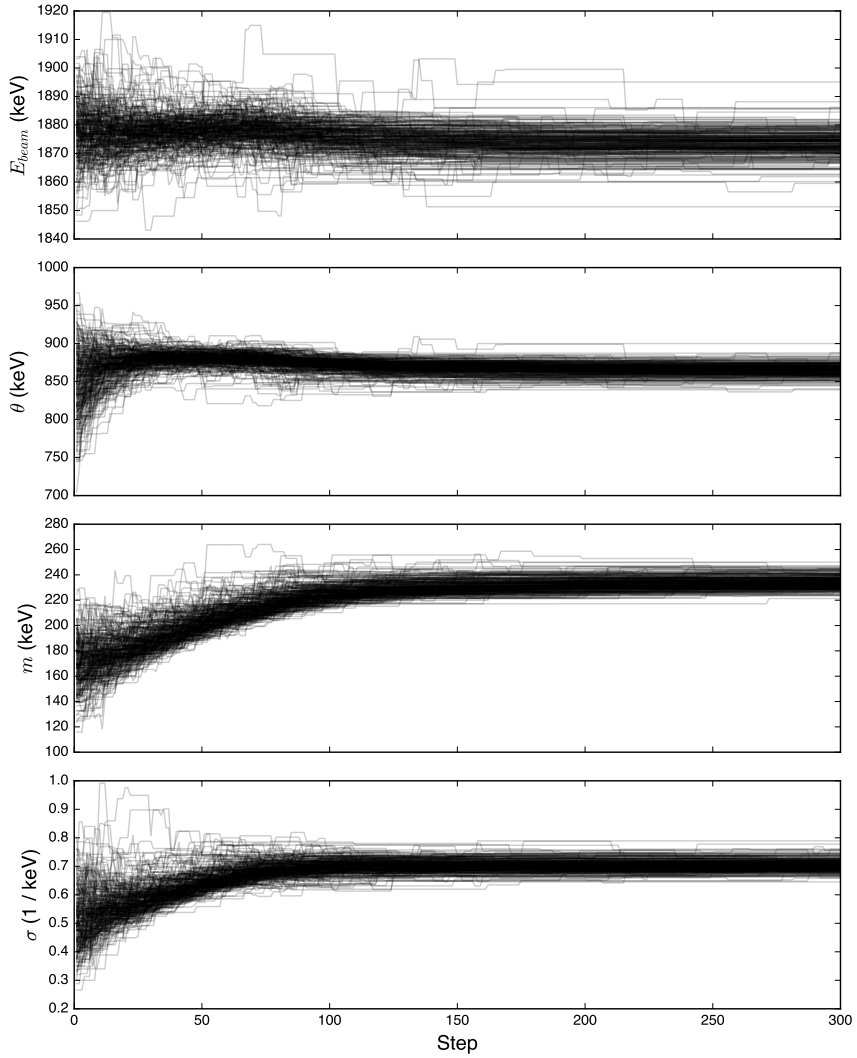

Figure D.3: Ensemble of chains from the MCMC fit to experimental TOF spectra showing the 4 parameters related to the description of the initial deuteron energy. Parameters are named according the conventions of Equations (D.2) and (D.3). For each parameter, the ensemble of 256 "walkers" is shown at each step in the chain. Evolution of the values is evident as the parameter space is explored and the ensemble moves towards higher likelihood values, away from the initial parameter estimates and narrowing around those which yield the highest likelihood for the observed TOF data using the assumed model.

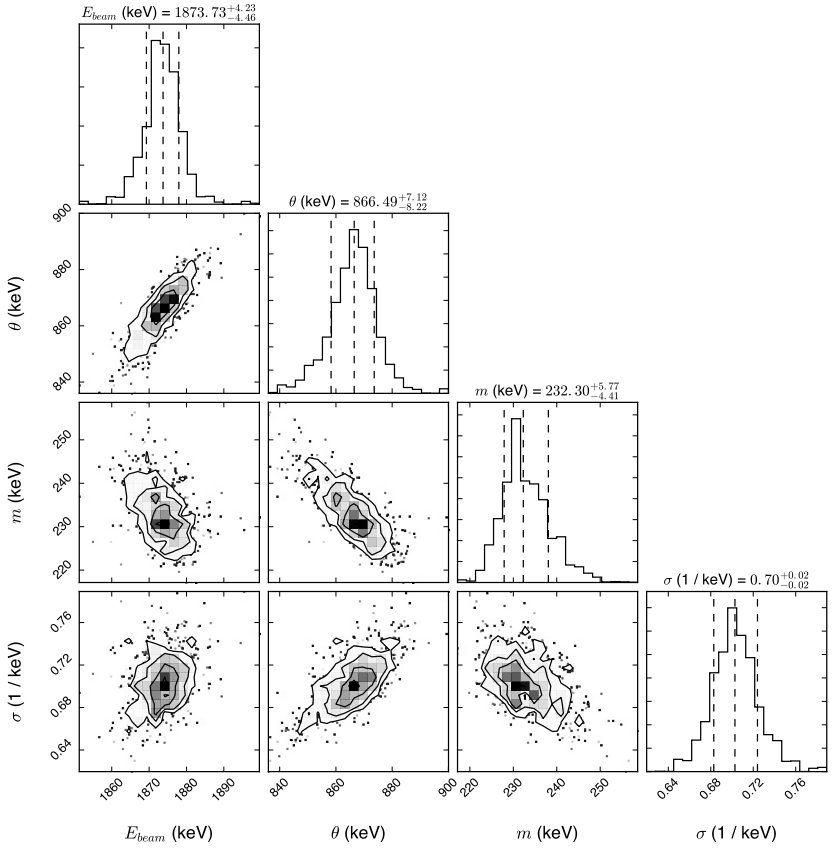

Figure D.4: "Corner plot" from MCMC result also shown in Figure D.3. Correlations between parameter values can be seen in most cases; these correlations arise completely within the fitting process and are not the result of any *a priori* knowledge.

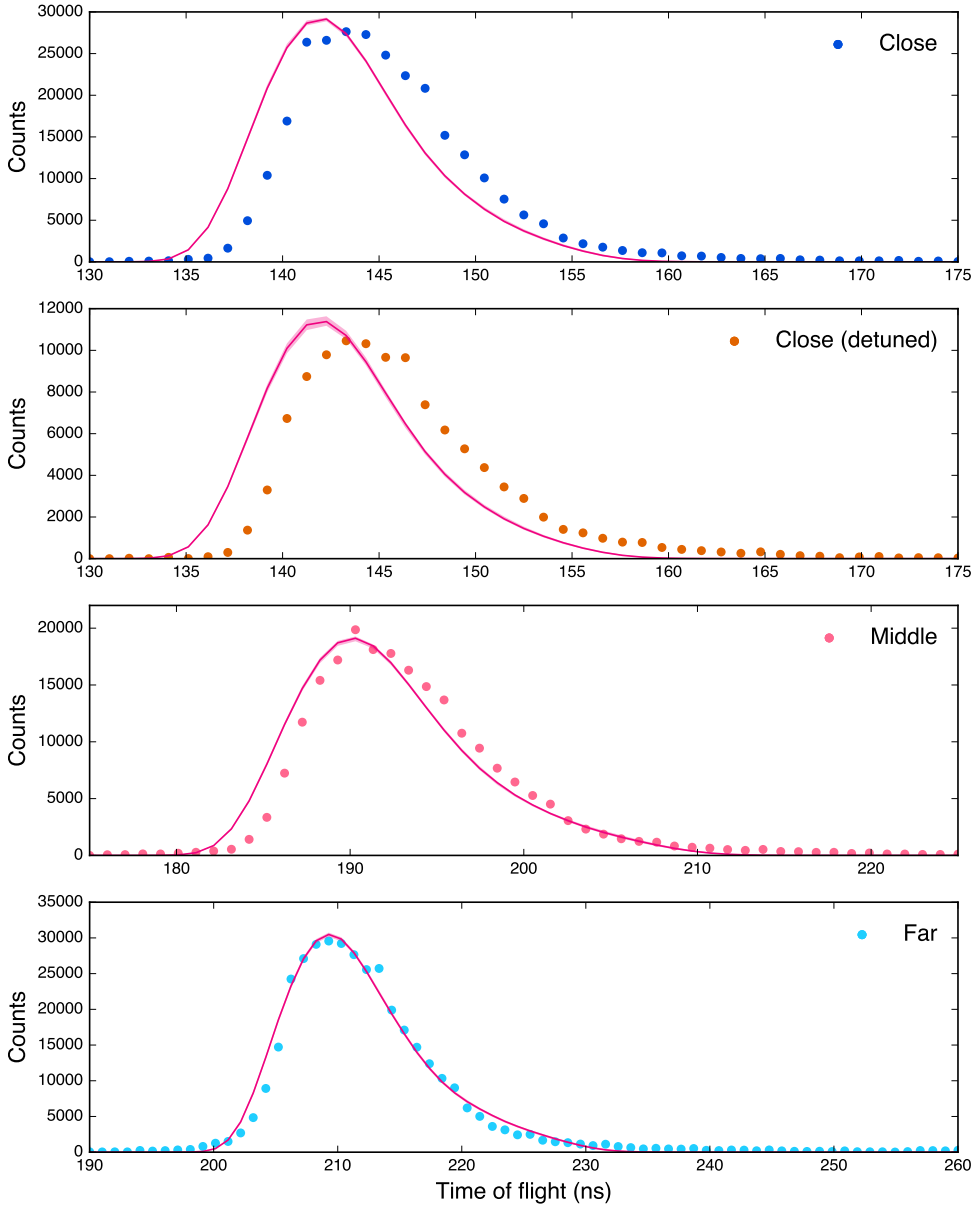

Figure D.5: Shown are experimental TOF spectra for 3 different standoff distances over 4 different data collection periods. The purpose of the "detuned" run is discussed in Section D.3). Overlaid on the TOF spectra is the model TOF from the MCMC fit to these data. Deviations from the observed TOF are evident, particularly in the shortest standoff runs, but these are insignificant compared to the width of the deuteron energy distribution, shown in Figure D.6.

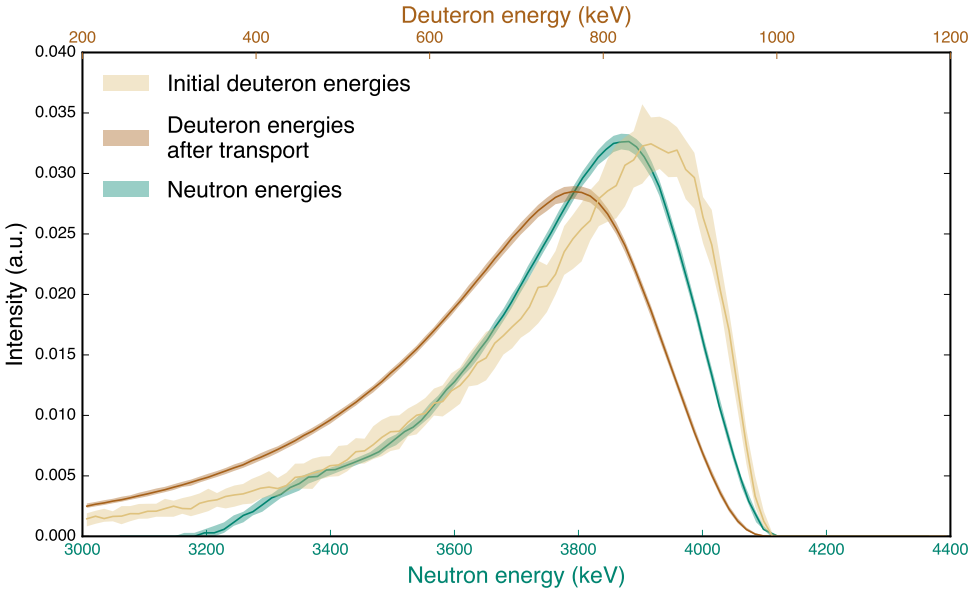

Figure D.6: The neutron and deuteron energy distributions from the MCMC fit, marginalized over the length of the gas cell. This is the neutron energy distribution used in the MCNPX-PoliMi simulations, as referred to in Section 4.4.1. The deuteron energy distribution is shown at both the beginning and the end of the gas cell, drawn as light beige and ochre, respectively.

APPENDIX E: SOURCE CODE

Section E.1: CMA filter implementation in C++ using ROOT and TWaveform

```
//
// this function produces the filtered and jump-rejecting baseline
// it populates the array passed into it as address rollingPwave
// assumes rollingPwave is already allocated and appropriate in size
// width sets the width of the averaging window
//
template<typename _Tp>
void getRollingPwave( TTemplWaveform<_Tp>* waveform, double* rollingPwave,
                      int halfWidth, double preloadValue,
                      double rejectThreshold ) {

    std::deque<double> movingBaselineFilter;
    baseline_t baselineDummy;
    baselineDummy.baseline = 0;
    baselineDummy.FWHM = 0;
    baselineDummy.chisquare = 0;



    double movingBaselineValue = 0;


    if( preloadValue > -7777 ) {
        /* check if preloadValue argument is a 'valid' sample value */
        /* default argument is -7777 */
        /* mildly arbitrary, but decidedly invalid for samples between 0 and 16383 */
        /* if a good preload value is supplied: */
        /*      then we treat the preload value as the "first half" of the of the boxcar at the start of the wave */
        /*      this helps stabilize the filter against the presence of signal in the first bit of the waveform */
        for( int i = 0; i < halfWidth; i++ ) {
            movingBaselineFilter.push_back( preloadValue );
            movingBaselineValue += preloadValue;
        }
    }
```

```
35
36      if( DEBUG >= 2 ) {
37          printf("after preload, deque has %lu entries\n", movingBaselineFilter.size() );
38          printf("average value is %.2f\n", movingBaselineValue / movingBaselineFilter.size()
        );
39      }
40
41
42
43      for( int i = 0; i < halfWidth; i++ ) {
44          if( preloadValue > -7777 ) {
45              if( fabs(getWaveformValueSafely( waveform, baselineDummy, i, false ) -
        movingBaselineValue/movingBaselineFilter.size() ) >= rejectThreshold ) {
46                  if( DEBUG >= 2 ) {
47                      printf("skipping early value %.2f at sample %i\n",
        getWaveformValueSafely( waveform, baselineDummy, i, false), i );
48                  }
49                  continue;
50              }
51          }
52          movingBaselineFilter.push_back( getWaveformValueSafely(waveform, baselineDummy, i,
        false ) );
53          movingBaselineValue += getWaveformValueSafely( waveform, baselineDummy, i, false );
54      }
55
56
57      for( int i = 0; i < waveform->GetLength(); i++ ) {
58
59          if( fabs(getWaveformValueSafely( waveform, baselineDummy, i+halfWidth, false ) -
        movingBaselineValue/movingBaselineFilter.size() ) < rejectThreshold  ) {
60              if( i + halfWidth < waveform->GetLength() ) {
61                  /* if we're here, we're still within valid lengths of the waveform */
62                  if( movingBaselineFilter.size() >= halfWidth*2 + 1 ) {
63                      /* here, the filter is fully occupied */
64                      /* so, pop values off the back as we go */
65                      movingBaselineValue -= movingBaselineFilter.front();
66                      movingBaselineFilter.pop_front();
67                  }
68
69                  movingBaselineValue += getWaveformValueSafely( waveform, baselineDummy, i+
        halfWidth, false );
```

```
70                  movingBaselineFilter.push_back(getWaveformValueSafely( waveform,
    baselineDummy, i+halfWidth, false ) );
71              }
72              else {
73                  movingBaselineValue -= movingBaselineFilter.front();
74                  movingBaselineFilter.pop_front();
75              }
76          }
77          rollingPwave[i] = movingBaselineValue / movingBaselineFilter.size();
78      }
79
80
81      if( DEBUG >= 2 ) {
82          printf("moving baseline deque size: %lu\n", movingBaselineFilter.size() );
83          printf("getRollingPwave - dumping first 10 values of waveform...\n");
84          for( int i = 0; i < 10; i++ ) {
85              printf("%.2f\t", rollingPwave[i] );
86          }
87          printf("\n");
88      }
89  }
```

chapter_sourceCode/cmaCode.cc

BIBLIOGRAPHY

[1] C.E. Aalseth et al. "Search for an Annual Modulation in a p-Type Point Contact Germanium Dark Matter Detector". *Phys. Rev. Lett.* 107.14 (2011). DOI: 10.1103/PhysRevLett.107.141301. arXiv: 1106.0650 [astro-ph.CO].

[2] R. Abbiati, E. Gatti, A. Geraci, and G. Ripamonti. "A new digital estimation technique for baseline restoration". *Nucl. Instrum. Meth. A* 548.3 (2005). DOI: 10.1016/j.nima.2005.04.072.

[3] J.N. Abdurashitov et al. "The Russian-American Gallium Experiment (SAGE) Cr Neutrino Source Measurement". *Phys. Rev. Lett.* 77.23 (1996). DOI: 10.1103/PhysRevLett.77.4708.

[4] S. Abrahamyan et al. "Measurement of the Neutron Radius of ^{208}Pb through Parity Violation in Electron Scattering". *Phys. Rev. Lett.* 108.11 (2012). DOI: 10.1103/PhysRevLett.108.112502. arXiv: 1201.2568 [nucl-ex].

[5] P. Agnes et al. "Results from the first use of low radioactivity argon in a dark matter search". *Phys. Rev. D* 93.8 (2016). DOI: 10.1103/PhysRevD.93.081101. arXiv: 1510.00702 [astro-ph.CO].

[6] G. Agnolet et al. "Background studies for the MINER Coherent Neutrino Scattering reactor experiment". *Nucl. Instrum. Meth. A* 853 (2017). DOI: 10.1016/j.nima.2017.02.024.

[7] M. Agostini et al. "Background-free search for neutrinoless double-β decay of ^{76}Ge with GERDA". *Nature* 544.7648 (2017). DOI: 10.1038/nature21717.

[8] A. Aguilar et al. "Evidence for neutrino oscillations from the observation of $\overline{\nu}_e$ appearance in a $\overline{\nu}_\mu$ beam". *Phys. Rev. D* 64.11 (2001). DOI: 10.1103/PhysRevD.64.112007.

[9] A.A. Aguilar-Arevalo et al. "Event Excess in the MiniBooNE Search for $\overline{\nu}_\mu \to \overline{\nu}_e$ Oscillations". *Phys. Rev. Lett.* 105.18 (2010). DOI: 10.1103/PhysRevLett.105.181801.

[10] S.P. Ahlen, F.T. Avignone, R.L. Brodzinski, A.K. Drukier, G. Gelmini, and D.N. Spergel. "Limits on cold dark matter candidates from an ultralow background germanium spectrometer". *Phys. Lett. B* 195.4 (1987). DOI: 10.1016/0370-2693(87)91581-4.

[11] S. Ahlen et al. "The case for a directional dark matter detector and the status of current experimental efforts". *Int. J. Mod. Phys. A* 25.01 (2010). DOI: 10.1142/S0217751X10048172. arXiv: 0911.0323 [astro-ph.CO].

[12] M. Akashi-Ronquest et al. "Improving Photoelectron Counting and Particle Identification in Scintillation Detectors with Bayesian Techniques". *Astropart. Phys.* 65 (2015). DOI: 10.1016/j.astropartphys.2014.12.006. arXiv: 1408.1914 [physics.ins-det].

[13] D. S. Akerib et al. "First Results from the LUX Dark Matter Experiment at the Sanford Underground Research Facility". *Phys. Rev. Lett.* 112.9 (2014). DOI: 10.1103/PhysRevLett.112.091303. arXiv: 1310.8214 [astro-ph.CO].

[14] D. Akimov et al. "The COHERENT experiment at the spallation neutron source" (2015). arXiv: 1509.08702 [physics.ins-det].

[15] D.Yu. Akimov et al. "Prospects for observation of neutrino-nuclear neutral current coherent scattering with two-phase Xenon emission detector". *J. Instrum.* 8.10 (2013). DOI: 10.1088/1748-0221/8/10/P10023. arXiv: 1212.1938 [physics.ins-det].

[16] D.Yu. Akimov et al. "RED-100 detector for the first observation of the elastic coherent neutrino scattering off xenon nuclei". In: *J. Phys. Conf. Ser.* Vol. 675. 1. 2016.

[17] D. Akimov et al. "Observation of coherent elastic neutrino-nucleus scattering". *Science* 357.6356 (2017). DOI: 10.1126/science.aao0990. arXiv: 1708.01294 [nucl-ex].

[18] B. Alex Brown. "Neutron Radii in Nuclei and the Neutron Equation of State". *Phys. Rev. Lett.* 85.25 (2000). DOI: 10.1103/PhysRevLett.85.5296.

[19] P.S. Amanik and G.C. McLaughlin. "Nuclear neutron form factor from neutrino–nucleus coherent elastic scattering". *J. Phys. G Nucl. Partic.* 36.1 (2008). DOI: 10.1088/0954-3899/36/1/015105.

[20] A.J. Anderson, J.M. Conrad, E. Figueroa-Feliciano, C. Ignarra, G. Karagiorgi, K. Scholberg, M.H. Shaevitz, and J. Spitz. "Measuring active-to-sterile neutrino oscillations with neutral current coherent neutrino-nucleus scattering". *Phys. Rev. D* 86.1 (2012). DOI: 10.1103/PhysRevD.86.013004. arXiv: 1201.3805 [hep-ph].

[21] P. Anselmann et al. "First results from the ^{51}Cr neutrino source experiment with the GALLEX detector". *Phys. Lett. B* 342.1-4 (1995). DOI: 10.1016/0370-2693(94)01586-2.

[22] E. Aprile et al. "Limits on Spin-Dependent WIMP-Nucleon Cross Sections from 225 Live Days of XENON100 Data". *Phys. Rev. Lett.* 111.2 (2013). DOI: 10.1103/PhysRevLett.111.021301. arXiv: 1301.6620 [astro-ph.CO].

[23] B. Armbruster et al. "Upper limits for neutrino oscillations $\overline{\nu}_\mu \to \overline{\nu}_e$ from muon decay at rest". *Phys. Rev. D* 65.11 (2002). DOI: 10.1103/PhysRevD.65.112001. arXiv: 0203021 [hep-ex].

[24] F. Arneodo et al. "Scintillation efficiency of nuclear recoil in liquid xenon". *Nucl. Instrum. Meth. A* 449.1 (2000). DOI: 10.1016/S0168-9002(99)01300-5.

[25] G. Arnison et al. "Experimental observation of lepton pairs of invariant mass around 95 GeV/c^2 at the CERN SPS collider". *Phys. Lett. B* 126.5 (1983). DOI: 10.1016/0370-2693(83)90188-0.

[26] R.G. Arns. "Detecting the neutrino". *Phys. Perspect.* 3.3 (2001). DOI: 10.1007/PL00000535.

[27] M. Auger et al. "Search for Neutrinoless Double-Beta Decay in ^{136}Xe with EXO-200". *Phys. Rev. Lett.* 109.3 (2012). DOI: 10.1103/PhysRevLett.109.032505. arXiv: 1205.5608 [hep-ex].

[28] Y.Y. Azmy, J.M. Barnes, J.D. Drischler, J.O. Johnston, R.A. Lillie, G.S. McNeilly, and R.T. Santoro. *Spallation Neutron Source Radiation Shielding Issues*. Tech. rep. Oak Ridge National Lab., TN (US), 1999. URL: https://www.osti.gov/scitech/servlets/purl/10411.

[29] Po. Bagnaia et al. "Evidence for $Z^0 \to e^+e^-$ at the CERN pp Collider". *Phys. Lett. B* 129.1-2 (1983). DOI: 10.1016/0370-2693(83)90744-X.

[30] J.N. Bahcall, A.M. Serenelli, and S. Basu. "New solar opacities, abundances, helioseismology, and neutrino fluxes". *Astrophys. J. Lett.* 621.1 (2005). DOI: 10.1086/428929. arXiv: 0412440 [astro-ph].

[31] H. Bai, Z. Wang, L. Zhang, Y. Lu, H. Jiang, J. Chen, and G. Zhang. "Calibration of an EJ309 liquid scintillator using an AmBe neutron source". *Nucl. Instrum. Meth. A* (2017). DOI: 10.1016/j.nima.2017.04.028.

[32] A.B. Balantekin and G.M. Fuller. "Neutrinos in cosmology and astrophysics". *Prog. Part. Nucl. Phys.* 71 (2013). DOI: 10.1016/j.ppnp.2013.03.008. arXiv: 1303.3874 [nucl-th].

[33] P.S. Barbeau, J.I. Collar, Yu. Efremenko, and K. Scholberg. "Comment on "Fitting the Annual Modulation in DAMA with Neutrons from Muons and Neutrinos"". *Phys. Rev. Lett.* 113.22 (2014). DOI: 10.1103/PhysRevLett.113.229001. arXiv: 1409.3185 [hep-ph].

[34] P.S. Barbeau, J.I. Collar, and O. Tench. "Large-mass ultralow noise germanium detectors: performance and applications in neutrino and astroparticle physics". *J. Cosmol. Astropart. P.* 2007.09 (2007). DOI: 10.1088/1475-7516/2007/09/009. arXiv: 0701012 [nucl-ex].

[35] R. Barlow and C. Beeston. "Fitting using finite Monte Carlo samples". *Comput. Phys. Commun.* 77.2 (1993). DOI: 10.1016/0010-4655(93)90005-W.

[36] J. Barranco, D. Delepine, M. Napsuciale, and A. Yebra. "New upper bound for the neutrino magnetic moment from its Dirac/Majorana nature and Borexino data" (2017). arXiv: 1704.01549 [hep-ph].

[37] J. Barranco, O.G. Miranda, and T.I. Rashba. "Probing new physics with coherent neutrino scattering off nuclei". *J. High Energy Phys.* 2005.12 (2005). DOI: 10.1088/1126-6708/2005/12/021.

[38] J. Barranco, O.G. Miranda, and T.I. Rashba. "Sensitivity of low energy neutrino experiments to physics beyond the standard model". *Phys. Rev. D* 76.7 (2007). DOI: 10.1103/PhysRevD.76.073008.

[39] K. Bays et al. "Supernova relic neutrino search at super-Kamiokande". *Phys. Rev. D* 85.5 (2012). DOI: 10.1103/PhysRevD.85.052007. arXiv: 1111.5031 [hep-ex].

[40] J.F. Beacom. "The diffuse supernova neutrino background". *Annu. Rev. Nucl. Part. S.* 60 (2010). DOI: 10.1146/annurev.nucl.010909.083331. arXiv: 1004.3311 [astro-ph.HE].

[41] A.G. Beda, V.B. Brudanin, V.G. Egorov, D.V. Medvedev, V.S. Pogosov, M.V. Shirchenko, and A.S. Starostin. "The results of search for the neutrino magnetic moment in GEMMA experiment". *Advances in High Energy Physics* (2012). DOI: 10.1155/2012/350150.

[42] O. Behnke, K. Kröninger, G. Schott, and T. Schörner-Sadenius, eds. *Data Analysis in High Energy Physics: A Practical Guide to Statistical Methods*. Weinheim: WILEY-VCH Verlag GmbH & Co. KGaA, 2013.

[43] M.J. Berger, J.S. Coursey, M.A. Zucker, and J. Chang. *ESTAR, PSTAR, and ASTAR: Computer Programs for Calculating Stopping-Power and Range Tables for Electrons, Protons, and Helium Ions*. Version 1.2.3. Gaithersburg, MD, 2005. URL: http://physics.nist.gov/Star (visited on 02/17/2017).

[44] J. Beringer et al. "Review of Particle Physics, 2012-2014". *Phys. Rev. D* 86.1 (2012). DOI: 10.1103/PhysRevD.86.010001.

[45] R. Bernabei et al. "First results from DAMA/LIBRA and the combined results with DAMA/NaI". *Eur. Phys. J. C* 56.3 (2008). DOI: 10.1140/epjc/s10052-008-0662-y.

[46] R. Bernabei et al. "No role for neutrons, muons and solar neutrinos in the DAMA annual modulation results". *Eur. Phys. J. C* 74.12 (2014).

[47] G. Bertone, D. Hooper, and J. Silk. "Particle dark matter: Evidence, candidates and constraints". *Phys. Rep.* 405.5 (2005). DOI: 10.1016/j.physrep.2004.08.031.

[48] H. Bethe and R. Peierls. "The "Neutrino"". *Nature* 133.3362 (1934).

[49] H.A. Bethe. "Supernova mechanisms". *Rev. Mod. Phys.* 62.4 (1990). DOI: 10.1103/RevModPhys.62.801.

[50] H.A. Bethe and J.R. Wilson. "Revival of a stalled supernova shock by neutrino heating". *Astrophys. J.* 295 (1985). DOI: 10.1086/163343.

[51] C. Bhatia et al. "Dual-fission chamber and neutron beam characterization for fission product yield measurements using monoenergetic neutrons". *Nucl. Instrum. Meth. A* 757 (2014). DOI: 10.1016/j.nima.2014.03.022.

[52] J.B. Birks. "Scintillations from organic crystals: specific fluorescence and relative response to different radiations". *P. Phys. Soc. A* 64.10 (1951). DOI: 10.1088/0370-1298/64/10/303.

[53] J.B. Birks. *The Theory and Practice of Scintillation Counting*. Vol. 27. International Series of Monographs in Electronics and Instrumentation. London: Pergamon, 1964.

[54] E. Bjorkland and D. Thompson. *SNS Timing Master Functional System Description*. Tech. rep. SNS 109020200-TD0006. Spallation Neutron Source, 2011.

[55] J.M. Blondin, A. Mezzacappa, and C. DeMarino. "Stability of standing accretion shocks, with an eye toward core-collapse supernovae". *Astrophys. J.* 584.2 (2003). DOI: 10.1086/345812. arXiv: 0210634 [astro-ph].

[56] W. Blum, W. Riegler, and L. Rolandi. *Particle detection with drift chambers*. 2nd ed. Berlin: Springer Science & Business Media, 2008.

[57] R. Bowers and J.R. Wilson. "Collapse of iron stellar cores". *Astrophys. J.* 263 (1982). DOI: 10.1086/160510.

[58] S.J. Brice et al. "A method for measuring coherent elastic neutrino-nucleus scattering at a far off-axis high-energy neutrino beam target". *Phys. Rev. D* 89.7 (2014). DOI: 10.1103/PhysRevD.89.072004. arXiv: 1311.5958 (physics.ins-det).

[59] E.M. Bringa and R.E. Johnson. "Coulomb Explosion and Thermal Spikes". *Phys. Rev. Lett.* 88.16 (2002). DOI: 10.1103/PhysRevLett.88.165501. arXiv: 0103475 [cond-mat].

[60] S.W. Bruenn. "Stellar core collapse-Numerical model and infall epoch". *Astrophys. J. Supplement Series* 58 (1985). DOI: 10.1086/191056.

[61] R. Brun and F. Rademakers. "ROOT—an object oriented data analysis framework". *Nucl. Instrum. Meth. A* 389.1-2 (1997). DOI: 10.1016/S0168-9002(97)00048-X.

[62] C. Buck, J. Hakenmüller, G. Heusser, M. Lindner, W. Maneschg, T. Rink, T. Schierhuber, H. Strecker, and V. Wagner. "CONUS: Towards the detection of coherent elastic neutrino nucleus scattering". In: *Proceedings of the 15th Conference on Topics in Astroparticle and Underground Physics*. (July 27, 2017). Sudbury, ON CA, 2017. URL: https://indico.cern.ch/event/606690/contributions/2591545.

[63] R. Buras, M. Rampp, H.-Th. Janka, and K. Kifonidis. "Improved Models of Stellar Core Collapse and Still No Explosions: What Is Missing?" *Phys. Rev. Lett.* 90.24 (2003). DOI: 10.1103/PhysRevLett.90.241101. arXiv: 0303171 [astro-ph].

[64] A. Burrows. "Colloquium: Perspectives on Core-Collapse Supernova Theory". *Rev. Mod. Phys.* 85.1 (2013). DOI: 10.1103/RevModPhys.85.245. arXiv: 1210.4921 [astro-ph.SR].

[65] A. Burrows, K. Klein, and R. Gandhi. "The future of supernova neutrino detection". *Phys. Rev. D* 45.10 (1992). DOI: 10.1103/PhysRevD.45.3361.

[66] B. Cabrera, L.M. Krauss, and F. Wilczek. "Bolometric detection of neutrinos". *Phys. Rev. Lett.* 55.1 (1985). DOI: 10.1103/PhysRevLett.55.25.

[67] Costruzioni Apparecchiature Elettroniche Nucleari S.p.A. (CAEN). *CAEN V1730 product page*. 2017. URL: http://www.caen.it/csite/CaenProd.jsp?idmod=779&parent=11 (visited on 08/23/2017).

[68] P. Cameron et al. "Spallation Neutron Source Ring Diagnostics". In: *Proceedings of the 2003 Particle Accelerator Conference*. (May 12–16, 2003). Portland, OR. URL: https://accelconf.web.cern.ch/accelconf/p03/PAPERS/WPPB013.PDF.

[69] H. Chagani, P. Majewski, E.J. Daw, V.A. Kudryavtsev, and N.J.C. Spooner. "Measurement of the quenching factor of Na recoils in NaI(Tl)". *J. Instrum.* 3.06 (2008). DOI: 10.1088/1748-0221/3/06/P06003.

[70] J.A. Christen and C. Fox. *A General Purpose Scale-Independent MCMC Algorithm*. Tech. rep. I-07-16. Guanajuato, Mexico: Centro de Investigación en Matemáticas (CIMAT), Dec. 12, 2007. URL: http://random.cimat.mx/reportes/enlinea/I-07-16.pdf.

[71] D.D. Cohen. "X-rays from an 241Am source and their relative intensities". *Nucl. Instrum. Meth. A* 267.2-3 (1988). DOI: 10.1016/0168-9002(88)90490-1.

[72] S.A. Colgate and R.H. White. "The hydrodynamic behavior of supernovae explosions". *Astrophys. J.* 143 (1966). DOI: 10.1086/148549.

[73] CDMS Collaboration. *CDMS Collaboration Website*. 2017. URL: http://cdms.berkeley.edu/.

[74] CRESST Collaboration. *CRESST Collaboration Website*. 2017. URL: http://www.cresst.de/.

[75] J. I. Collar. "Quenching and channeling of nuclear recoils in NaI(Tl): Implications for dark-matter searches". *Phys. Rev. C* 88.3 (2013). DOI: 10.1103/PhysRevC.88.035806. arXiv: 1302.0796 [physics.ins-det].

[76] J.I. Collar. "Light WIMP searches: the effect of the uncertainty in recoil energy scale and quenching factor". *arXiv preprint arXiv:1010.5187* (2010).

[77] J.I. Collar, N.E. Fields, M. Hai, T.W. Hossbach, J.L. Orrell, C.T. Overman, G. Perumpilly, and B. Scholz. "Coherent neutrino-nucleus scattering detection with a CsI [Na] scintillator at the SNS spallation source". *Nucl. Instrum. Meth. A* 773 (2015). DOI: 10.1016/j.nima.2014.11.037.

[78] C.L. Cowan, F. Reines, F.B. Harrison, H.W. Kruse, and A.D. McGuire. "Detection of the Free Neutrino: a Confirmation". *Science* 124.3212 (1956). ISSN: 0036-8075. DOI: 10.1126/science.124.3212.103. eprint: http://science.sciencemag.org/content/124/3212/103.full.pdf.

[79] A.S. Crowell. "Cross-Section Measurements of Star Configurations in Neutron-Deuteron Breakup at 16.0 MeV". PhD thesis. Duke University, 2001.

[80] St. Gobain Crystals. *BC-501, BC-501A, BC-519 products page*. 2017. URL: http://www.crystals.saint-gobain.com/products/bc501-bc501a-bc519 (visited on 08/22/2017).

[81] St. Gobain Crystals. *CsI[Tl] and CsI[Na] Material Data Sheet*. 2016. URL: http://www.crystals.saint-gobain.com/sites/imdf.crystals.com/files/documents/csitl-and-na-material-data-sheet_69771.pdf (visited on 09/15/2017).

[82] P. Cushman et al. "Snowmass CF1 Summary: WIMP Dark Matter Direct Detection" (2013). arXiv: 1310.8327 [hep-ex].

[83] B. Dasgupta, A. Mirizzi, I. Tamborra, and R. Tomàs. "Neutrino mass hierarchy and three-flavor spectral splits of supernova neutrinos". *Phys. Rev. D* 81.9 (2010). DOI: 10.1103/PhysRevD.81.093008. arXiv: 0802.1481 [hep-ph].

[84] J.H. Davis. "Fitting the Annual Modulation in DAMA with Neutrons from Muons and Neutrinos". *Phys. Rev. Lett.* 113.8 (2014). DOI: 10.1103/PhysRevLett.113.081302. arXiv: 1407.1052 [hep-ph].

[85] H. Davoudiasl, H.-S. Lee, and W.J. Marciano. "Muon Anomaly and Dark Parity Violation". *Phys. Rev. Lett.* 109.3 (2012). DOI: 10.1103/PhysRevLett.109.031802. arXiv: 1205.2709 [hep-ph].

[86] H. Davoudiasl, H.-S. Lee, and W.J. Marciano. "Muon $g-2$, rare kaon decays, and parity violation from dark bosons". *Phys. Rev. D* 89.9 (2014). DOI: 10.1103/PhysRevD.89.095006. arXiv: 1402.3620 [hep-ph].

[87] H. De Vries, C.W. De Jager, and C. De Vries. "Nuclear charge-density-distribution parameters from elastic electron scattering". *Atom. Data Nucl. Data* 36.3 (1987). DOI: 10.1016/0092-640X(87)90013-1.

[88] A. Deandrea. "Atomic parity violation in cesium and implications for new physics". *Phys. Lett. B* 409.1 (1997). arXiv: 9705435 [hep-ph].

[89] P. Denzel, J. Diemand, and R. Angélil. "Molecular dynamics simulations of bubble nucleation in dark matter detectors". *Phys. Rev. E* 93.1 (2016). DOI: 10.1103/PhysRevE.93.013301. arXiv: 1601.07390 [physics.comp-ph].

[90] R. Diener, S. Godfrey, and I. Turan. "Constraining extra neutral gauge bosons with atomic parity violation measurements". *Phys. Rev. D* 86.11 (2012). DOI: 10.1103/PhysRevD.86.115017. arXiv: 1111.4566 [hep-ph].

[91] ORNL Neutron Sciences Directorate. *ORNL SNS website*. 2017. URL: https://neutrons.ornl.gov/sns (visited on 07/06/2017).

[92] D.L. Donoho and I.M. Johnstone. "Adapting to Unknown Smoothness via Wavelet Shrinkage". *J. Am. Stat. Assoc.* 90.432 (1995). DOI: 10.1080/01621459.1995.10476626.

[93] R. Dossi, A. Ianni, G. Ranucci, and O.Ju. Smirnov. "Methods for precise photoelectron counting with photomultipliers". *Nucl. Instrum. Meth. A* 451.3 (2000). DOI: 10.1016/S0168-9002(00)00337-5.

[94] A. Drukier and L. Stodolsky. "Principles and applications of a neutral-current detector for neutrino physics and astronomy". *Phys. Rev. D* 30.11 (1984). DOI: 10.1103/PhysRevD.30.2295.

[95] H. Duan, G.M. Fuller, J. Carlson, and Y.-Z. Qian. "Neutrino Mass Hierarchy and Stepwise Spectral Swapping of Supernova Neutrino Flavors". *Phys. Rev. Lett.* 99.24 (2007). DOI: 10.1103/PhysRevLett.99.241802. arXiv: 0707.0290 [astro-ph].

[96] C.A. Duba, F. Duncan, J. Farine, A. Habig, A. Hime, et al. "HALO: The helium and lead observatory for supernova neutrinos". *J. Phys. Conf. Ser.* 136 (2008). DOI: 10.1088/1742-6596/136/4/042077.

[97] B. Dutta, Y. Gao, A. Kubik, R. Mahapatra, N. Mirabolfathi, L.E. Strigari, and J.W. Walker. "Sensitivity to oscillation with a sterile fourth generation neutrino from ultralow threshold neutrino-

nucleus coherent scattering". *Phys. Rev. D* 94.9 (2016). DOI: 10.1103/PhysRevD.94.093002. arXiv: 1511.02834 [hep-ph].

[98] B. Dutta, R. Mahapatra, L.E. Strigari, and J.W. Walker. "Sensitivity to Z-prime and nonstandard neutrino interactions from ultralow threshold neutrino-nucleus coherent scattering". *Phys. Rev. D* 93.1 (2016). DOI: 10.1103/PhysRevD.93.013015. arXiv: 1508.07981 [hep-ph].

[99] Yu. Efremenko and W.R. Hix. "Opportunities for neutrino physics at the Spallation Neutron Source (SNS)". In: *Proceedings of Carolina International Symposium on Neutrino Physics*. (May 15–17, 2008). Vol. 173. 1. Columbia, SC. DOI: 10.1088/1742-6596/173/1/012006.

[100] S.R. Elliott. "Measuring supernova neutrino temperatures using lead perchlorate". *Phys. Rev. C* 62.6 (2000). DOI: 10.1103/PhysRevC.62.065802. arXiv: 0006041 [astro-ph].

[101] J. Engel. "Nuclear form factors for the scattering of weakly interacting massive particles". *Phys. Lett. B* 264.1-2 (1991). DOI: 10.1016/0370-2693(91)90712-Y.

[102] J. Engel, G.C. McLaughlin, and C. Volpe. "What can be learned with a lead-based supernova-neutrino detector?" *Phys. Rev. D* 67.1 (2003). DOI: 10.1103/PhysRevD.67.013005. arXiv: 0209267 [hep-ph].

[103] A. Enqvist, C.C. Lawrence, B.M. Wieger, S.A. Pozzi, and T.N. Massey. "Neutron light output response and resolution functions in EJ-309 liquid scintillation detectors". *Nucl. Instrum. Meth. A* 715 (2013). DOI: 10.1016/j.nima.2013.03.032.

[104] G. Fernandez Moroni, J. Estrada, E.E. Paolini, G. Cancelo, J. Tiffenberg, and J. Molina. "Charge coupled devices for detection of coherent neutrino-nucleus scattering". *Phys. Rev. D* 91.7 (2015). DOI: 10.1103/PhysRevD.91.072001.

[105] R.L. Fleischer, P.B. Price, and R.M. Walker. "Ion explosion spike mechanism for formation of charged-particle tracks in solids". *J. Appl. Phys* 36.11 (1965). DOI: 10.1063/1.1703059.

[106] D. Foreman-Mackey. "corner.py: Scatterplot matrices in Python". *The Journal of Open Source Software* 24 (2016). DOI: 10.21105/joss.00024.

[107] D. Foreman-Mackey et al. *emcee*. https://github.com/dfm/emcee. 2010.

[108] D. Foreman-Mackey, D.W. Hogg, D. Lang, and J. Goodman. "emcee: The MCMC hammer". *Publ. Astron. Soc. Pac.* 125.925 (2013). DOI: 10.1086/670067. arXiv: 1202.3665 [astro-ph.IM].

[109] J.A. Formaggio, E. Figueroa-Feliciano, and A.J. Anderson. "Sterile neutrinos, coherent scattering, and oscillometry measurements with low-temperature bolometers". *Phys. Rev. D* 85.1 (2012). DOI: 10.1103/PhysRevD.85.013009. arXiv: 1107.3512 [hep-ph].

[110] J.A. Formaggio and C.J. Martoff. "Backgrounds to sensitive experiments underground". *Annu. Rev. Nucl. Part. S.* 54 (2004). DOI: 10.1146/annurev.nucl.54.070103.181248.

[111] D.Z. Freedman. "Coherent effects of a weak neutral current". *Phys. Rev. D* 9.5 (1974). DOI: 10.1103/PhysRevD.9.1389.

[112] D.Z. Freedman, D.N. Schramm, and D.L. Tubbs. "The weak neutral current and its effects in stellar collapse". *Annu. Rev. Nucl. Sci.* 27.1 (1977). DOI: 10.1146/annurev.ns.27.120177.001123.

[113] K. Freese, J. Frieman, and A. Gould. "Signal modulation in cold-dark-matter detection". *Phys. Rev. D* 37.12 (1988). DOI: 10.1103/PhysRevD.37.3388.

[114] G.M. Fuller, W.C. Haxton, and G.C. McLaughlin. "Prospects for detecting supernova neutrino flavor oscillations". *Phys. Rev. D* 59.8 (1999). DOI: 10.1103/PhysRevD.59.085005. arXiv: 9809164 [astro-ph].

[115] Th.K. Gaisser and M. Honda. "Flux of atmospheric neutrinos". *Annu. Rev. Nucl. Part. S.* 52.1 (2002). DOI: 10.1146/annurev.nucl.52.050102.090645. arXiv: 0203272 [hep-ph].

[116] R.J. Gaitskell. "Direct detection of dark matter". *Annu. Rev. Nucl. Part. S.* 54 (2004). DOI: 10.1146/annurev.nucl.54.070103.181244.

[117] S. Geman and D. Geman. "Stochastic Relaxation, Gibbs Distributions, and the Bayesian Restoration of Images". *IEEE T. Pattern. Anal.* 6 (1984). DOI: 10.1109/TPAMI.1984.4767596.

[118] G. Gerbier et al. "Pulse shape discrimination and dark matter search with NaI(Tl) scintillator". *Astropart. Phys.* 11.3 (1999). DOI: 10.1088/1748-0221/3/06/P06003.

[119] W. R. Gibbs and J.-P. Dedonder. "Neutron radii of the calcium isotopes". *Phys. Rev. C* 46.5 (1992). DOI: 10.1103/PhysRevC.46.1825.

[120] G.K. Giovanetti et al. "A dark matter search with MALBEK". *Phys. Proc.* 61 (2015). DOI: 10.1016/j.phpro.2014.12.014.

[121] S.L. Glashow. "Partial-symmetries of weak interactions". *Nucl. Phys.* 22.4 (1961). ISSN: 0029-5582. DOI: 10.1016/0029-5582(61)90469-2.

[122] M.C. Gonzalez-Garcia, M. Maltoni, and T. Schwetz. "Global analyses of neutrino oscillation experiments". *Nucl. Phys. B* 908 (2016). DOI: 10.1016/j.nuclphysb.2016.02.033.

[123] J. Goodman and J. Weare. "Ensemble Samplers with Affine Invariance". *Comm. App. Math. And Comp. Sci.* 5.1 (2010). DOI: 10.2140/camcos.2010.5.65.

[124] M.W. Goodman and E. Witten. "Detectability of certain dark-matter candidates". *Phys. Rev. D* 31.12 (1985). DOI: 10.1103/PhysRevD.31.3059.

[125] D. Groom. *Atomic and nuclear properties of deuterium gas (D2).* 2016. URL: http://pdg.lbl.gov/2016/AtomicNuclearProperties/HTML/deuterium_gas.html (visited on 01/29/2017).

[126] C. Guo et al. "Neutron beam tests of CsI (Na) and CaF 2 (Eu) crystals for dark matter direct search". *Nucl. Instrum. Meth. A* 818 (2016). DOI: 10.1016/j.nima.2016.02.037.

[127] K.B.W. Harpsøe, M.I. Andersen, and P. Kjægaard. "Bayesian photon counting with electron-multiplying charge coupled devices (EMCCDs)". *Astron. Astrophys.* 537 (2012). DOI: 10.1051/0004-6361/201117089.

[128] F.J. Hasert et al. "Observation of neutrino-like interactions without muon or electron in the Gargamelle neutrino experiment". *Phys. Lett. B* 46.1 (1973). DOI: 10.1016/0370-2693(73)90499-1.

[129] F.J. Hasert et al. "Search for elastic muon-neutrino electron scattering". *Phys. Lett. B* 46.1 (1973). ISSN: 0370-2693. DOI: 10.1016/0370-2693(73)90494-2.

[130] R.H. Helm. "Inelastic and Elastic Scattering of 187-MeV Electrons from Selected Even-Even Nuclei". *Phys. Rev.* 104.5 (1956). DOI: 10.1103/PhysRev.104.1466.

[131] R.G. Herb. "The Pelletron Accelerator". *IEEE T. Nucl. Sci.* 18.3 (1971). DOI: 10.1109/TNS.1971.4325962.

[132] G. Heusser. "Low-radioactivity background techniques". *Annu. Rev. Nucl. Part. S.* 45.1 (1995). DOI: 10.1146/annurev.ns.45.120195.002551.

[133] G. Heyes, W.A. Watson III, J.C. Jastrzembski, D. Abbott, and D. Barker. "The CEBAF on-line data acquisition System". In: *Proceedings of the 1994 International Conference on Computing in High Energy and Nuclear Physics (CHEP)*. (Apr. 21–27, 1994). San Francisco, CA USA, 1994.

[134] R. Hofstadter. "Nuclear and Nucleon Scattering of High-Energy Electrons". *Annu. Rev. Nucl. Sci.* 7.1 (1957). DOI: 10.1146/annurev.ns.07.120157.001311. eprint: https://doi.org/10.1146/annurev.ns.07.120157.001311.

[135] C.J. Horowitz. "Neutrino trapping in a supernova and the screening of weak neutral currents". *Phys. Rev. D* 55.8 (1997). DOI: 10.1103/PhysRevD.55.4577. arXiv: 9603138 [astro-ph].

[136] C.J. Horowitz, K.J. Coakley, and D.N. McKinsey. "Supernova observation via neutrino-nucleus elastic scattering in the CLEAN detector". *Phys. Rev. D* 68.2 (2003). DOI: 10.1103/PhysRevD.68.023005.

[137] C.J. Horowitz, K.S. Kumar, and R. Michaels. "Electroweak measurements of neutron densities in CREX and PREX at JLab, USA". *Eur. Phys. J. A* 50.2 (2014). DOI: 10.1140/epja/i2014-14048-3. arXiv: 1307.3572 [nucl-ex].

[138] C.J. Horowitz and J. Piekarewicz. "Neutron Star Structure and the Neutron Radius of ^{208}Pb". *Phys. Rev. Lett.* 86.25 (2001). DOI: 10.1103/PhysRevLett.86.5647. arXiv: 0010227 [astro-ph].

[139] C.J. Horowitz, S.J. Pollock, P.A. Souder, and R. Michaels. "Parity violating measurements of neutron densities". *Phys. Rev. C* 63.2 (2001). DOI: 10.1103/PhysRevC.63.025501. arXiv: 9912038 [nucl-th].

[140] P.Q. Hung and J.J. Sakurai. "The structure of neutral currents". *Annu. Rev. Nucl. Part. S.* 31.1 (1981). DOI: 10.1146/annurev.ns.31.120181.002111.

[141] C. Iliadis. *Nuclear Physics of Stars*. 1st ed. Weinheim: WILEY-VCH Verlag GmbH & Co. KGaA, 2007.

[142] Proteus Inc. *Proteus company website*. 2016. URL: http://www.proteus-pp.com.

[143] N. Itoh, R. Asahara, N. Tomizawa, S. Wanajo, and S. Nozawa. "Ion-ion correlation effect on the neutrino-nucleus scattering in supernova cores". *Astrophys. J.* 611.2 (2004). DOI: 10.1086/422380. arXiv: 0401488 [astro-ph].

[144] J. Iwanowska, L. Swiderski, T. Krakowski, M. Moszynski, T. Szczesniak, and G. Pausch. "The time-of-flight method for characterizing the neutron response of liquid organic scintillators". *Nucl. Instrum. Meth. A* 781 (2015). DOI: 10.1016/j.nima.2015.01.051.

[145] F. James and M. Roos. "Minuit: A System for Function Minimization and Analysis of the Parameter Errors and Correlations". *Comput. Phys. Commun.* 10 (1975). DOI: 10.1016/0010-4655(75)90039-9.

[146] H.-T. Janka. "Explosion Mechanisms of Core-Collapse Supernovae". *Annu. Rev. Nucl. Part. S.* 62 (2012). DOI: 10.1146/annurev-nucl-102711-094901. arXiv: 1206.2503 [astro-ph.SR].

[147] T.H.Y. Joshi. "A novel source of tagged low-energy nuclear recoils". *Nucl. Instrum. Meth. A* 656.1 (2011). DOI: 10.1016/j.nima.2011.07.044.

[148] T.H. Joshi et al. "First Measurement of the Ionization Yield of Nuclear Recoils in Liquid Argon". *Phys. Rev. Lett.* 112.17 (2014). DOI: 10.1103/PhysRevLett.112.171303. arXiv: 1402.2037 [physics.ins-det].

[149] G. Jungman, M. Kamionkowski, and K. Griest. "Supersymmetric dark matter". *Phys. Rep.* 267.5-6 (1996). DOI: 10.1016/0370-1573(95)00058-5.

[150] B. Kayser. "Majorana neutrinos and their electromagnetic properties". *Phys. Rev. D* 26.7 (1982). DOI: 10.1103/PhysRevD.26.1662.

[151] M. Kesselman and W.C. Dawson. "Spallation Neutron Source Beam Current Monitor Electronics". In: *Presented at* 10^{th} *Beam Instrumentation Workshop*. (May 6–9, 2002). Brookhaven National Lab.

[152] S.R. Klein and J. Nystrand. "Exclusive vector meson production in relativistic heavy ion collisions". *Phys. Rev. C* 60.1 (1999). DOI: 10.1103/PhysRevC.60.014903. arXiv: 9902259 [hep-ph].

[153] U Kneissl, HH Pitz, and A Zilges. "Investigation of nuclear structure by resonance fluorescence scattering". *Prog. Part. Nucl. Phys.* 37 (1996). DOI: 10.1016/0146-6410(96)00055-5.

[154] G.F. Knoll. *Radiation detection and measurement*. 4th ed. Hoboken, NJ: John Wiley & Sons, 2010.

[155] E. Kolbe and K. Langanke. "Role of ν – induced reactions on lead and iron in neutrino detectors". *Phys. Rev. C* 63.2 (2001). DOI: 10.1103/PhysRevC.63.025802. arXiv: 0003060 [nucl-th].

[156] A. Konovalov. PhD thesis. Moscow Engineering Physics Intsitute, 2019.

[157] J. Kopp, P.A.N. Machado, M. Maltoni, and T. Schwetz. "Sterile neutrino oscillations: the global picture". *J. High Energy Phys.* 2013.5 (2013). DOI: 10.1007/JHEP05(2013)050. arXiv: 1303.3011 [hep-ph].

[158] *Spallation Neutron Source Final Safety Assessment Document For Neutron Facilities*. Tech. rep. SNS 102030102-ES0016-R03. Spallation Neutron Source, 2011. URL: https://conference.sns.gov/event/28/attachments/65/84/FSAD-Neutron_Facilities_Revision_3.pdf.

[159] N.V. Kornilov, I. Fabry, S. Oberstedt, and F.-J. Hambsch. "Total characterization of neutron detectors with a 252 Cf source and a new light output determination". *Nucl. Instrum. Meth. A* 599.2 (2009). DOI: 10.1016/j.nima.2008.10.032.

[160] T.S. Kosmas, O.G. Miranda, D.K. Papoulias, M. Tórtola, and J.W.F. Valle. "Probing neutrino magnetic moments at the Spallation Neutron Source facility". *Phys. Rev. D* 92.1 (2015). DOI: 10.1103/PhysRevD.92.013011. arXiv: 1505.03202 [hep-ph].

[161] T.S. Kosmas, D.K. Papoulias, M. Tórtola, and J.W.F. Valle. "Probing light sterile neutrino signatures at reactor and Spallation Neutron Source neutrino experiments" (2017). arXiv: 1703.00054 [hep-ph].

[162] K.S. Krane. *Introductory Nuclear Physics*. 3rd ed. Vol. 465. New York: Wiley, 1987.

[163] J.M. Lattimer and M. Prakash. "Nuclear matter and its role in supernovae, neutron stars and compact object binary mergers". *Phys. Rep.* 333 (2000). arXiv: 0002203 [astro-ph].

[164] J.M. Lattimer and M. Prakash. "The Physics of Neutron Stars". *Science* 304.5670 (2004). ISSN: 0036-8075. DOI: 10.1126/science.1090720. arXiv: 0405262 [astro-ph].

[165] J.A. LaVerne, A. Hitachi, and T. Doke. "Scintillation response of liquid argon for 3–20 MeV He ions". *Nucl. Instrum. Meth. A* 327.1 (1993). DOI: 10.1016/0168-9002(93)91411-F.

[166] G.P. Lawrence, R.K. Beauchamp, and J.L. McKibben. "Direct extraction of negative ion beams of good intensity from a duoplasmatron". *Nucl. Instrum. Methods* 32.2 (1965). DOI: 10.1016/0029-554X(65)90539-2.

[167] A.F. Leder et al. "Measuring Neutron Spectrum at MIT Research Reactor Utilizing He-3 Bonner Cylinder Approach with an Unfolding Analysis" (2017). arXiv: 1710.00802 [physics.ins-det].

[168] E.J. Lentz, A. Mezzacappa, O.E.B. Messer, W.R. Hix, and S.W. Bruenn. "Interplay of neutrino opacities in core-collapse supernova simulations". *Astrophys. J.* 760.1 (2012). DOI: 10.1088/0004-637X/760/1/94.

[169] W.R. Leo. *Techniques for nuclear and particle physics experiments: a how-to approach*. Springer Science & Business Media, 1987.

[170] J.D. Lewin and P.F. Smith. "Review of mathematics, numerical factors, and corrections for dark matter experiments based on elastic nuclear recoil". *Astropart. Phys.* 6.1 (1996).

[171] J. Liao and D. Marfatia. "COHERENT constraints on nonstandard neutrino interactions" (2017). arXiv: 1708.04255 [hep-ph].

[172] H. Liskien and A. Paulsen. "Neutron production cross sections and energies for the reactions ^{7}Li$(p,n)\,^{7}$Be and ^{7}Li$(p,n)\,^{7}$Be*". *Atom. Data Nucl. Data* 15.1 (1975). DOI: 10.1016/0092-640X(75)90004-2.

[173] H. Liskien and A. Paulsen. "Neutron production cross sections and energies for the reactions T$(p,n)\,^{3}$He, D$(d,n)\,^{3}$He, and T$(d,n)\,^{4}$He". *Atom. Data Nucl. Data* 11.7 (1973).

[174] C. Lunardini and A.Y. Smirnov. "Probing the neutrino mass hierarchy and the 13-mixing with supernovae". *J. Cosmol. Astropart. P.* 2003.06 (2003). DOI: 10.1088/1475-7516/2003/06/009. arXiv: 0302033 [hep-ph].

[175] T. Lund, A. Marek, C. Lunardini, H.-T. Janka, and G. Raffelt. "Fast time variations of supernova neutrino fluxes and their detectability". *Phys. Rev. D* 82.6 (2010). DOI: 10.1103/PhysRevD.82.063007. arXiv: 1006.1889 [astro-ph.HE].

[176] A. Manzur, A. Curioni, L. Kastens, D. N. McKinsey, K. Ni, and T. Wongjirad. "Scintillation efficiency and ionization yield of liquid xenon for monoenergetic nuclear recoils down to 4 keV". *Phys. Rev. C* 81.2 (2010). DOI: 10.1103/PhysRevC.81.025808.

[177] A. Marek, H.-Th. Janka, R. Buras, M. Liebendörfer, and M. Rampp. "On ion-ion correlation effects during stellar core collapse". *Astron. Astrophys.* 443.1 (2005). DOI: 10.1051/0004-6361:20053236. arXiv: 0504291 [astro-ph].

[178] M.G. Marino. *TWaveform*. https://github.com/mgmarino/TWaveform. 2012.

[179] T.E. Mason et al. "The Spallation Neutron Source in Oak Ridge: A powerful tool for materials research". *Physica B* 385 (2006). DOI: j.physb.2006.05.281.

[180] R.J. McConn, C.J. Gesh, R.T. Pagh, R.A. Rucker, and R. Williams III. *Compendium of material composition data for radiation transport modeling*. Tech. rep. PNNL-15870 Rev. 1. Richland, WA, USA: Pacific Northwest National Laboratory (PNNL), Mar. 4, 2011. URL: http://www.pnnl.gov/main/publications/external/technical_reports/PNNL-15870Rev1.pdf.

[181] A. Meldrum, S.J. Zinkle, L.A. Boatner, and R.C. Ewing. "A transient liquid-like phase in the displacement cascades of zircon, hafnon and thorite". *Nature* 395.6697 (1998). DOI: 10.1038/25698.

[182] G. Mention, M. Fechner, Th. Lasserre, Th. A. Mueller, D. Lhuillier, M. Cribier, and A. Letourneau. "Reactor antineutrino anomaly". *Phys. Rev. D* 83.7 (2011). DOI: 10.1103/PhysRevD.83.073006. arXiv: 1101.2755 [hep-ex].

[183] Mesytec. *MPD-4 datasheet*. Version 2.12_01. 2015. URL: http://www.mesytec.com/products/datasheets/MPD-4.pdf (visited on 08/02/2016).

[184] N. Metropolis, A.W. Rosenbluth, M.N. Rosenbluth, A.H. Teller, and E. Teller. "Equation of State Calculations by Fast Computing Machines". *J. Chem. Phys.* 21.6 (1953). DOI: 10.1063/1.1699114.

[185] J.P. Miller, E. De Rafael, and B.L. Roberts. "Muon $(g-2)$: experiment and theory". *Rep. Prog. Phys.* 70.5 (2007). arXiv: 0703049 [hep-ph].

[186] O. Moreno and T. W. Donnelly. "Unified approach to electron and neutrino elastic scattering off nuclei with an application to the study of the axial structure". *Phys. Rev. C* 92.5 (2015). DOI: 10.1103/PhysRevC.92.055504. arXiv: 1506.04733 [nucl-th].

[187] National Instruments. *PCI-5153 Oscilloscope Device support page*. 2017. URL: http://www.ni.com/en-us/support/model.pci-5153.html (visited on 07/14/2017).

[188] R.M. Neal. *Probabilistic Inference Using Markov Chain Monte Carlo Methods*. Tech. rep. CRG-TR-93-1. Toronto, Ontario, Canada: Department of Computer Science, University of Toronto, Sept. 25, 1993.

[189] NIST/SEMATECH. "Lognormal Distribution". In: *e-Handbook of Statistical Methods*. Ed. by C. Croarkin, P. Tobias, and J.J. Filliben. Oct. 30, 2013 ed. 2003. URL: http://www.itl.nist.gov/div898/handbook/eda/section3/eda3669.htm (visited on 08/23/2017).

[190] M.A. Norsworthy, A. Poitrasson-Rivière, M.L. Ruch, S.D. Clarke, and S.A. Pozzi. "Evaluation of neutron light output response functions in EJ-309 organic scintillators". *Nucl. Instrum. Meth. A* 842 (2017). DOI: 10.1016/j.nima.2016.10.035.

[191] D.R. Nygren. "Columnar recombination: a tool for nuclear recoil directional sensitivity in a xenon-based direct detection WIMP search". In: *J. Phys. Conf. Ser.* Vol. 460. 1. 2013. DOI: 10.1088/1742-6596/460/1/012006.

[192] Kaye & Laby Online. "4.2.1 X-ray absorption edges, characteristic X-ray lines and fluorescence yields". In: *Tables of Physical & Chemical Constants (16^{th} edition 1995)*. Version 1.0. 2005. URL: http://www.kayelaby.npl.co.uk/atomic_and_nuclear_physics/4_2/4_2_1.html.

[193] N. Paar, D. Vretenar, T. Marketin, and P. Ring. "Inclusive charged-current neutrino-nucleus reactions calculated with the relativistic quasiparticle random phase approximation" (2007). arXiv: nucl-th/0710.4881.

[194] J. Papavassiliou, J. Bernabéu, and M. Passera. "Neutrino-nuclear coherent scattering and the effective neutrino charge radius". In: *Proceedings of the International Europhysics Conference on High Energy Physics,* (July 21–27, 2005). Lisbon, Portugal, 2005. arXiv: 0512029 [hep-ph].

[195] H. Park et al. "Neutron beam test of CsI crystal for dark matter search". *Nucl. Instrum. Meth. A* 491.3 (2002). DOI: 10.1016/S0168-9002(02)01274-3.

[196] C. Patrignani et al. "Review of Particle Physics, 2016-2017". *Chinese Phys. C* 40 (2016). DOI: 10.1088/1674-1137/40/10/100001.

[197] K. Patton, J. Engel, G.C. McLaughlin, and N. Schunck. "Neutrino-nucleus coherent scattering as a probe of neutron density distributions". *Phys. Rev. C* 86.2 (2012). DOI: 10.1103/PhysRevC.86.024612. arXiv: 1207.0693 [nucl-th].

[198] D.G. Phillips II et al. "The MAJORANA experiment: an ultra-low background search for neutrinoless double-beta decay". In: *J. Phys. Conf. Ser.* Vol. 381. 1. 2012. DOI: 10.1088/1742-6596/381/1/012044. arXiv: 1111.5578 [nucl-ex].

[199] *Photomultiplier Tubes and Assemblies for Scintillation Counting and High Energy Physics*. Hamamatsu Photonics K.K. Sept. 2012.

[200] F. Pino, L. Stevanato, D. Cester, G. Nebbia, L. Sajo-Bohus, and G. Viesti. "The light output and the detection efficiency of the liquid scintillator EJ-309". *Appl. Radiat. Isotopes* 89 (2014). DOI: 10.1016/j.apradiso.2014.02.016.

[201] S.J. Pollock and M.C. Welliver. "Effects of neutron spatial distributions on atomic parity nonconservation in cesium". *Phys. Lett. B* 464.3 (1999). DOI: 10.1016/S0370-2693(99)00987-9. arXiv: 9904062 [nucl-th].

[202] S. G. Porsev, K. Beloy, and A. Derevianko. "Precision Determination of Electroweak Coupling from Atomic Parity Violation and Implications for Particle Physics". *Phys. Rev. Lett.* 102.18 (2009). DOI: 10.1103/PhysRevLett.102.181601. arXiv: 0902.0335 [hep-ph].

[203] S. G. Porsev, K. Beloy, and A. Derevianko. "Precision determination of weak charge of ^{133}Cs from atomic parity violation". *Phys. Rev. D* 82.3 (2010). DOI: 10.1103/PhysRevD.82.036008. arXiv: 1006.4193 [hep-ph].

[204] S.A. Pozzi et al. "MCNPX-PoliMi for nuclear nonproliferation applications". *Nucl. Instrum. Meth. A* 694 (2012). DOI: 10.1016/j.nima.2012.07.040.

[205] J.R. Prescott. "A statistical model for photomultiplier single-electron statistics". *Nucl. Instrum. Meth.* 39.1 (1966). DOI: 10.1016/0029-554X(66)90059-0.

[206] B. Pritychenko and A. Sonzogni. *Q-value Calculator (QCalc)*. 2017. URL: http://www.nndc.bnl.gov/qcalc/ (visited on 09/11/2017).

[207] Y.-Z. Qian, W.C. Haxton, K. Langanke, and P. Vogel. "Neutrino-induced neutron spallation and supernova *r*-process nucleosynthesis". *Phys. Rev. C* 55.3 (1997). DOI: 10.1103/PhysRevC.55.1532. arXiv: 9611010 [nucl-th].

[208] *R11265U Series / H11934 Series data sheet*. Hamamatsu Photonics K.K. Mar. 2013.

[209] P.-G. Reinhard and W. Nazarewicz. "Information content of a new observable: The case of the nuclear neutron skin". *Phys. Rev. C* 81.5 (2010). DOI: 10.1103/PhysRevC.81.051303. arXiv: 1002.4140 [nucl-th].

[210] M. Rennich and V. Graves. "Design and Operational Features of a Mercury Target Facility". In: *Presented at Neutrino Factory International Design Study Kick-off Meeting*. (Dec. 15–17, 2008). CERN.

[211] G.C. Rich. *The lead neutrino cube*. Tech. rep. COHERENT Collaboration, 2016.

[212] G.C. Rich, K. Kazkaz, H.P. Martinez, and T. Gushue. "Fabrication and characterization of a lithium-glass-based composite neutron detector". *Nucl. Instrum. Meth. A* 794 (2015). DOI: 10.1016/j.nima.2015.05.004.

[213] D. Rimal, M. McIntyre, and H. Ray. *SNS Neutrino Flux Simulation: Technical Note*. Tech. rep. COHERENT Collaboration, 2015.

[214] W.A. Rolke, A.M. Lopez, and J. Conrad. "Limits and confidence intervals in the presence of nuisance parameters". *Nucl. Instrum. Meth. A* 551.2 (2005). DOI: 10.1016/j.nima.2005.05.068. arXiv: 0403059 [physics.data-an].

[215] T. Diaz de la Rubia, R.S. Averback, R. Benedek, and W.E. King. "Role of thermal spikes in energetic displacement cascades". *Phys. Rev. Lett.* 59.17 (1987). DOI: 10.1103/PhysRevLett.59.1930.

[216] V.C. Rubin and W.K. Ford Jr. "Rotation of the andromeda nebula from a spectroscopic survey of emission regions". *Astrophys. J.* 159 (1970). DOI: 10.1086/150317.

[217] T. Saab and E. Figueroa-Feliciano. *WIMP Limit Plotter*. Version 2.22. 2017. URL: http://cdms.berkeley.edu/limitplots/mm/WIMP_limit_plotter.html (visited on 02/09/2017).

[218] A. Salam and J.C. Ward. "Electromagnetic and weak interactions". *Phys. Lett.* 13.2 (1964). ISSN: 0031-9163. DOI: 10.1016/0031-9163(64)90711-5.

[219] S. Sangiorgio et al. "R&D for the Observation of Coherent Neutrino-Nucleus Scatter at a Nuclear Reactor with a Dual-Phase Argon Ionization Detector". *Phys. Proc.* 37 (2012). DOI: 10.1016/j.phpro.2012.02.465.

[220] D. Santos, F. Mayet, O. Guillaudin, Th. Lamy, S. Ranchon, A. Trichet, P. Colas, and I. Giomataris. "Ionization Quenching Factor Measurement of ^{4}He" (2008). arXiv: astro-ph/0810.1137.

[221] K. Scholberg. "Prospects for measuring coherent neutrino-nucleus elastic scattering at a stopped-pion neutrino source". *Phys. Rev. D* 73.3 (2006). DOI: 10.1103/PhysRevD.73.033005. arXiv: 0511042 [hep-ex].

[222] K. Scholberg, T. Wongjirad, E. Hungerford, A. Empl, D. Markoff, P. Mueller, Y. Efremenko, D. McKinsey, and J. Nikkel. "The CLEAR Experiment". In: *Proceedings of the DPF-2009 Conference*. (July 27–31, 2009). Detroit, MI USA, 2009. arXiv: 0910.1989 (physics.hep-ex).

[223] H. Schölermann and H. Klein. "Optimizing the energy resolution of scintillation counters at high energies". *Nucl. Instrum. Methods* 169.1 (1980). DOI: 10.1016/0029-554X(80)90097-X.

[224] B.J. Scholz. PhD thesis. University of Chicago, 2017.

[225] B.J. Scholz. *CsI[Na] Note 01*. Tech. rep. COHERENT Collaboration, 2015.

[226] D.N. Schramm and W.D. Arnett. "Neutral currents and supernovas". *Phys. Rev. Lett.* 34.2 (1975). DOI: 10.1103/PhysRevLett.34.113.

[227] Phillips Scientific. *6954 Single Channel GHz Bipolar Amplifier Product Page*. URL: http://www.phillipsscientific.com/preview/6954pre.htm (visited on 09/16/2017).

[228] Phillips Scientific. *Quad DC-300 MHz Amplifier family data sheet*. 1996. URL: http://www.phillipsscientific.com/pdf/770ds.pdf (visited on 09/29/2017).

[229] L.M. Sehgal. "Differences in the coherent interactions of ν_e, ν_μ, and ν_τ". *Phys. Lett. B* 162.4-6 (1985). DOI: 10.1016/0370-2693(85)90942-6.

[230] F. Seitz. "On the theory of the bubble chamber". *The Physics of Fluids* 1.1 (1958).

[231] S.J. Sheather. "Density Estimation". *Statistical Science* 19.4 (2004). DOI: 10.1214/088342304000000297.

[232] E. Simon et al. "SICANE: a detector array for the measurement of nuclear recoil quenching factors using a monoenergetic neutron beam". *Nucl. Instrum. Meth. A* 507.3 (2003). DOI: 10.1016/S0168-9002(03)01438-4.

[233] D. Sivia and J. Skilling. *Data Analysis: A Bayesian Tutorial*. Oxford, UK: Oxford University Press, 2006.

[234] N.J.C. Spooner, G.J. Davies, J.D. Davies, G.J. Pyle, T.D. Bucknell, G.T.A. Squier, J.D. Lewin, and P.F. Smith. "The scintillation efficiency of sodium and iodine recoils in a NaI(Tl) detector for dark matter searches". *Phys. Lett. B* 321.1 (1994). DOI: 10.1016/0370-2693(94)90343-3.

[235] T. Stiegler, C. Sofka, R.C. Webb, and J.T. White. "A study of the NaI(Tl) detector response to low energy nuclear recoils and a measurement of the quenching factor in NaI(Tl)" (2017). arXiv: 1706.07494 [physics.ins-det].

[236] R.J. Stokey and P.J. Lee. "Approximation to the Probability Density at the Output of a Photomultiplier Tube". In: *The Telecommunications and Data Acquisition Progress Report 42-73, January–March 1983*. Pasadena, CA: Jet Propulsion Laboratory, 1983. URL: http://ipnpr.jpl.nasa.gov/progress_report/42-73/73D.PDF.

[237] D.A. Swartz, A. Clocchiattit, R. Benjamin, D.F. Lester, and J.C. Wheeler. "Supernova 1993J as a spectroscopic link between type II and type Ib supernovae". *Nature* 365.6443 (1993). DOI: 10.1038/365232a0.

[238] I. Tamborra, F. Hanke, B. Müller, H.-T. Janka, and G. Raffelt. "Neutrino Signature of Supernova Hydrodynamical Instabilities in Three Dimensions". *Phys. Rev. Lett.* 111.12 (2013). DOI: 10.1103/PhysRevLett.111.121104. arXiv: 1307.7936 [astro-ph.SR].

[239] A. Tomanin, J. Paepen, P. Schillebeeckx, R. Wynants, R. Nolte, and A. Lavietes. "Characterization of a cubic EJ-309 liquid scintillator detector". *Nucl. Instrum. Meth. A* 756 (2014). DOI: 10.1016/j.nima.2014.03.028.

[240] D.R. Tovey, V. Kudryavtsev, M. Lehner, J.E. McMillan, C.D. Peak, J.W. Roberts, N.J.C. Spooner, and J.D. Lewin. "Measurement of scintillation efficiencies and pulse-shapes for nuclear recoils in NaI(Tl) and CaF_2(Eu) at low energies for dark matter experiments". *Phys. Lett. B* 433.1 (1998). DOI: 10.1016/S0370-2693(98)00643-1.

[241] J. Towns et al. "XSEDE: Accelerating Scientific Discovery". *Comput. Sci. Eng.* 16.5 (2014). ISSN: 1521-9615. DOI: 10.1109/MCSE.2014.80.

[242] V.I. Tretyak. "Semi-empirical calculation of quenching factors for ions in scintillators". *Astropart. Phys.* 33.1 (2010). DOI: 10.1016/j.astropartphys.2009.11.002. arXiv: 0911.3041 [nucl-ex].

[243] M.B. Tsang et al. "Constraints on the symmetry energy and neutron skins from experiments and theory". *Phys. Rev. C* 86.1 (2012). DOI: 10.1103/PhysRevC.86.015803. arXiv: 1204.0466 [nucl-ex].

[244] D. Väänänen and C. Volpe. "The neutrino signal at HALO: learning about the primary supernova neutrino fluxes and neutrino properties". *J. Cosmol. Astropart. P.* 2011 (2011). DOI: 10.1088/1475-7516/2011/10/019. arXiv: 1105.6225 [astro-ph.SR].

[245] D. Vale, T. Rauscher, and N. Paar. "Hybrid method to resolve the neutrino mass hierarchy by supernova (anti) neutrino induced reactions". *J. Cosmol. Astropart. P.* 2016.02 (2016). DOI: 10.1088/1475-7516/2016/02/007. arXiv: 1110.0103 [astro-ph].

[246] R.J. Van de Graaff. "Tandem electrostatic accelerators". *Nucl. Instrum. Methods* 8.2 (1960). DOI: 10.1016/S0029-554X(60)80006-7.

[247] W. Verkerke and D. Kirkby. "The RooFit toolkit for data modeling". In: *Proceedings of the 2003 International Conference on Computing in High Energy and Nuclear Physics (CHEP)*. (Mar. 24–28, 2003). La Jolla, CA US, 2003. arXiv: 0306116 [physics.data-an].

[248] M. Wang, G. Audi, F.G. Kondev, W.J. Huang, S. Naimi, and X. Xu. "The AME2016 atomic mass evaluation (II). Tables, graphs and references". *Chinese Phys. C* 41.3 (2017). DOI: 10.1088/1674-1137/41/3/030003.

[249] S. Weinberg. "A Model of Leptons". *Phys. Rev. Lett.* 19.21 (1967). DOI: 10.1103/PhysRevLett.19.1264.

[250] S. Weinberg. "Conceptual foundations of the unified theory of weak and electromagnetic interactions". *Rev. Mod. Phys.* 52.3 (1980). DOI: 10.1103/RevModPhys.52.515.

[251] C.R. Westerfeldt, E.G. Bilpuch, M.E. Bleck, D.A. Outlaw, W.K. Wells, J.F. Wilkerson, and T.B. Clegg. "A system for producing high-resolution polarized and unpolarized beams with a tandem accelerator". *Nucl. Instrum. Methods* 219.2 (1984). DOI: 10.1016/0167-5087(84)90334-X.

[252] J.F. Wilkerson, C. Torre, and R. Fauber. *TUNL FIELD utility*. June 1984. URL: http://www.tunl.duke.edu/field.php (visited on 01/30/2016).

[253] H.T. Wong et al. "Search of neutrino magnetic moments with a high-purity germanium detector at the Kuo-Sheng nuclear power station". *Phys. Rev. D* 75.1 (2007). DOI: 10.1103/PhysRevD.75.012001. arXiv: 0605006 [hep-ex].

[254] C.S. Wood, S.C. Bennett, D. Cho, B.P. Masterson, J.L. Roberts, C.E. Tanner, and C.E. Wieman. "Measurement of parity nonconservation and an anapole moment in cesium". *Science* 275.5307 (1997). DOI: 10.1126/science.275.5307.1759.

[255] S.E. Woosley, D.H. Hartmann, R.D. Hoffman, and W.C. Haxton. "The ν-process". *Astrophys. J.* 356 (1990).

[256] S.E. Woosley, A. Heger, and T. A. Weaver. "The evolution and explosion of massive stars". *Rev. Mod. Phys.* 74.4 (2002). DOI: 10.1103/RevModPhys.74.1015.

[257] J. Xu et al. "Scintillation efficiency measurement of Na recoils in NaI(Tl) below the DAMA/LIBRA energy threshold". *Phys. Rev. C* 92.1 (2015). DOI: 10.1103/PhysRevC.92.015807. arXiv: 1503.07212 [physics.ins-det].

[258] J.F. Ziegler, M.D. Ziegler, and J.P. Biersack. "SRIM–The stopping and range of ions in matter (2010)". *Nucl. Instrum. Meth. B* 268.11 (2010). DOI: 10.1016/j.nimb.2010.02.091.

[259] F. Zwicky. "Die Rotverschiebung von extragalaktischen Nebeln". *Helv. Phys. Acta* 6 (1933).

[260] F. Zwicky. "On the Masses of Nebulae and of Clusters of Nebulae". *Astrophys. J.* 86 (1937). DOI: 10.1086/143864.