

Performance of the Eos detector with water

S. Arora,^{1,2} M. Askins,^{1,2} A. J. Bacon,³ Z. Bagdasarian,^{1,2} A. Baldoni,^{4,5} L. Bartoszek,⁶ M. Bergevin,⁷ Y. Bezawada,^{1,2} E. Blucher,⁸ J. Boissevain,⁶ R. Bonventre,² E. J. Callaghan,^{1,2} D. F. Cowen,⁴ K. DeHolton,⁴ M. Diwan,⁹ M. Dubnowski,³ P. Englezos,^{1,2} S. Gadamssetty,^{1,2} C. Grant,¹⁰ B. Harris,³ M. R. Hebert,^{1,2} S. Jeon,¹⁰ T. Kaptanoglu,^{1,2,*} A. Katt,¹⁰ J. R. Klein,³ T. Kroupová,³ L. Lebanowski,^{1,2} S. Lynch,^{1,2} A. Mastbaum,¹¹ C. Mauger,³ G. Mayers,³ M. Miller,¹² J. Nachtman,¹² S. Naugle,³ J. Newby,¹³ M. Newcomer,³ A. Nikolica,³ G. D. Orebi Gann,^{1,2} A. Phipps,¹⁴ L. Pickard,^{1,2} R. C. Pitelka,³ L. Ren,¹⁵ A. Rincon,^{1,2} R. Rosero,⁹ N. Rowe,^{1,2} H. J. Ryoo,^{1,2} J. Ryshkewitch,¹⁰ J. Saba,² S. Schoppmann,^{1,2} J. Shen,³ M. Smiley,^{1,2} H. Song,^{1,2,10} H. Steiger,^{16,17} B. Tam,¹⁸ E. Tiras,^{19,12} W. H. To,²⁰ M. R. Vagins,²¹ R. Van Berg,⁶ J. Wallig,² G. Wendel,⁴ M. Wetstein,²² M. Wurm,¹⁶ G. Yang,^{1,2} M. Yeh,⁹ E. D. Zimmerman,¹⁵ and A. Zummo¹¹

¹Physics Department, *University of California at Berkeley*, Berkeley, California 94720-7300, USA

²*Lawrence Berkeley National Laboratory*, 1 Cyclotron Road, Berkeley, California 94720-8153, USA

³Department of Physics and Astronomy, *University of Pennsylvania*, Philadelphia, Pennsylvania 19104-6396, USA

⁴Department of Physics, *Pennsylvania State University*, University Park, Pennsylvania 16802, USA

⁵*United States Military Academy*, Department of Physics and Nuclear Engineering, West Point, New York 10996, USA

⁶Bartoszek Engineering, Aurora, Illinois 60506, USA

⁷*Lawrence Livermore National Laboratory*, Livermore, California 94550, USA

⁸*The Enrico Fermi Institute and Department of Physics, The University of Chicago*, Chicago, Illinois 60637, USA

⁹*Brookhaven National Laboratory*, Upton, New York 11973, USA

¹⁰*Boston University*, Department of Physics, Boston, Massachusetts 02215, USA

¹¹Department of Physics and Astronomy, *Rutgers, The State University of New Jersey*, 136 Frelinghuysen Road, Piscataway, New Jersey 08854-8019 USA

¹²Department of Physics and Astronomy, *The University of Iowa*, Iowa City, Iowa 52242, USA

¹³*Oak Ridge National Laboratory*, 1 Bethel Valley Road, Oak Ridge, Tennessee 37830

¹⁴*California State University, East Bay*, Department of Physics, 25800 Carlos Bee Boulevard, Hayward, California 94542, USA

¹⁵*University of Colorado at Boulder*, Department of Physics, Boulder, Colorado 80309, USA

¹⁶*Johannes Gutenberg-Universität*, Institute of Physics and EC PRISMA⁺⁺, Mainz, 55099 Mainz, Germany

¹⁷Physics Department, *Technische Universität München*, 85748 Garching, Germany

¹⁸Department of Physics, *University of Oxford*, Parks Road, Oxford OX1 3PU, United Kingdom

¹⁹Department of Physics, *Erciyes University*, 38030, Kayseri, Turkey

²⁰*California State University, Stanislaus*, Department of Physics, Turlock, California 95382, USA

²¹Department of Physics and Astronomy, *University of California, Irvine*, Irvine, California 92697-4575, USA

²²Department of Physics and Astronomy, *Iowa State University*, Ames, Iowa 50011, USA



(Received 15 June 2026; accepted 31 August 2026; published 21 September 2026)

In this manuscript we present the first results from Eos, a four tonne optical detector located at the University of California, Berkeley. The primary goal of Eos is to demonstrate the performance capabilities of scintillation-based, “hybrid” detector technology for future neutrino detectors. The data presented were collected while both the inner target vessel and the outer buffer vessel were filled with water. The water target acts as a well-understood medium that produces only Cherenkov light, which can be used to calibrate and develop the detector model and reconstruction algorithms prior to the deployment of scintillating material. Using deployed optical and radioactive calibration sources, a series of detailed detector calibrations are performed. These enable a suite of tests for various reconstruction algorithms. Simulations that use calibrated models are compared with the data across a variety of different types of calibration sources, source positions, and rotations.

DOI: [10.1103/7qkt-mnqz](https://doi.org/10.1103/7qkt-mnqz)

*Contact author: tannerbk@berkeley.edu

I. INTRODUCTION

The study of neutrinos is a tremendously broad field of science, with insights into the fusion processes in our sun, the explosion mechanism of distant supernovae, monitoring of terrestrial nuclear activity, and potentially the matter-antimatter asymmetry in the early Universe. The detection of neutrinos is a rich field, and historically significant progress has been made utilizing deep-underground detectors, such as Super-Kamiokande and SNO [1,2], filled with water. In these experiments, neutrino interactions produce Cherenkov light, which can be utilized to reconstruct the direction of the incoming neutrino. In more modern experiments, such as Borexino and SNO+ [3,4], liquid scintillator is chosen as the target material, which produces far more light than water and can be more easily cleaned of the radioactive contaminants that produce backgrounds. However, detecting the relatively small amount of Cherenkov light in a liquid scintillator detector is challenging, and has only been achieved statistically across many events [5] or at relatively high energy [6], in large-scale neutrino detectors.

Future detectors such as Theia [7] hope to utilize groundbreaking technology in order to simultaneously detect both Cherenkov and scintillation photons. This hybrid technique would allow for low thresholds, excellent energy and vertex resolution, event-by-event directionality, and particle identification that leverages the ratio of Cherenkov to scintillation photons. Theia offers world-leading sensitivity to a broad range of physics, including measurements of neutrinoless double beta decay, low-energy solar neutrinos, and the diffuse supernova neutrino background [8]. Additionally, hybrid detectors are being investigated for far-field detection of antineutrinos for the purpose of nuclear nonproliferation [9,10].

There are several technology developments that are being explored for hybrid event detection. Novel scintillators can emphasize the Cherenkov signal, either by reducing the scintillation light yield [11] or by slowing the scintillation emission timing [12–14]. Red-sensitive and/or fast photodetectors [15,16] can enhance the Cherenkov signal, particularly in combination with spectral photon sorting provided by devices like the dichroicon [17,18]. Advances in fast-electronics, waveform processing algorithms, and sophisticated reconstruction algorithms [19] will aid in the identification of the Cherenkov photons. Other ideas include unique approaches to photon collection to isolate the Cherenkov component in liquid argon neutrino detectors [20].

The Eos detector, described in detail in Sec. II, is a multiton-scale optical detector that utilizes many of these advanced technologies. The primary goal of Eos is to demonstrate the performance capabilities of the hybrid detector technology, and to test Monte Carlo (MC) simulation model predictions. This includes a demonstration of full event reconstruction, ultimately utilizing both

Cherenkov and scintillation light. These objectives largely drove the detector design and the selection of hardware and electronics for Eos. In the future, Eos plans to probe the dependence of the reconstruction resolution on detector configurations, such as the properties of the target material, the photocoverage, photon detector time resolution, and the addition of dichroicons for spectral sorting. MC models will be calibrated and validated against data, which will support predictions for next-generation experiments. More details about the detector goals are provided in Ref. [21].

Eos operated with the inner vessel filled with water in various time periods between May 2024 and April 2026 (the full water data collection period is referred to as the water phase). During this period, data were collected with deployed optical and radioactive calibration sources to evaluate the performance of the detector prior to introducing scintillation light, to calibrate the behavior of the various detector components, and to develop the infrastructure and analysis tools necessary for achieving the experimental goals. The relatively simple optical behavior of the water target makes this period critical for evaluating the baseline performance of Eos and provides a benchmark to which future phases can be compared. The future phases involve the deployment of both water-based liquid scintillator and “pure” liquid scintillator (e.g., LAB + PPO) as the primary target material. In this manuscript, we describe the detector in Sec. II, outline the simulation, processing, data collection, analysis tools, and event selection in Secs. III–VII respectively, present the full calibration of the detector in Sec. VIII, and finally detail the systematics, model verification, and water phase reconstruction performance results in Secs. IX–XI.

II. THE EOS DETECTOR

The Eos detector, shown in Fig. 1, has a four-tonne capacity acrylic inner vessel (IV), which can contain either a water or scintillator target. The IV is roughly cylindrical in shape, with ellipsoidal end caps and a 91 cm radius at the center. The IV is viewed by 241 PMTs, which detect the light created from interactions within the detector. Eos utilizes three different types of PMTs to study the impact of varying photon detector configurations. 204 Hamamatsu R14688-100 PMTs [22] are placed around the barrel and at the bottom of the detector. These eight-inch diameter PMTs have excellent single photoelectron (SPE) timing resolution, less than one ns FWHM, and high quantum efficiencies ($> 25\%$ at 400 nm) [23]. Twenty-four Hamamatsu twelve-inch R11780 PMTs [24] are placed at the top of the detector to increase the total light collection. An array of twelve dichroicons [18] are placed in front of twelve of the eight-inch R14688-100 PMTs at the bottom of the detector to provide spectral sorting of the Cherenkov and scintillation light [18]. Thirteen Hamamatsu ten-inch R7081 PMTs [25] are placed 35 cm behind the dichroicons to collect the light that passes through the dichroic filters. The PMTs and




FIG. 1. A schematic of the Eos detector showing the various detector components.

IV are contained within a 30-ton capacity stainless steel outer vessel (OV) filled with water. A neck that runs from the IV to the top lid of the detector allows the calibration sources to be deployed down the central axis. Figure 2




FIG. 2. A picture of the Eos detector taken during the construction, while the IV was filled air and prior to moving the detector components into the OV. This image is looking down through the acrylic vessel toward the eight-inch R14688-100 PMTs along the barrel as well as some of the dichroicons at the bottom of the detector. The twelve-inch R11780 PMTs above the acrylic vessel are seen at the top. The top surface and neck of the inner acrylic vessel are also visible. More images of the detector are shown in Ref. [26].

shows a picture of the detector, taken during construction, that shows many of the components, including the various PMT types, the dichroicons, and the IV. The detector is surrounded by a muon veto system consisting of optically isolated plastic scintillator panels coupled to PMTs. For the data collected in this manuscript, the muon veto system was not yet operational. Reference [26] provides more details about the detector design, configuration, and construction of the Eos detector.

The size of the Eos outer vessel was designed with practical limitations in mind, such as the size of the space available at UC Berkeley, the weight limitation of the overhead crane, and the width of the door into the space. With these limitations, the size of the detector was maximized, such that timing-based reconstruction algorithms could be tested. The choice to utilize the 12-inch and 10-inch PMTs (all repurposed from previous experiments) was made to maximize photocathode coverage, including coverage behind the dichroicons. The size of the inner vessel was limited by the space needed to fit the PMTs and, at the bottom, the dichroicons as well. The shape of the inner vessel was based on structural and practical needs, in coordination with the supplier, Reynolds Polymer Technology. The R14688 PMTs outperform the other PMTs, and 20 additional R14688 PMTs have been purchased since Eos has begun operating. These additional PMTs could be used to replace some or all of the other types of PMTs in the future. More details about how detector design choices were made is provided in Ref. [26].

Based on the mechanical tolerances of the detector and calibration deployment system as well as measurements performed while constructing Eos, the vertical position (z) of the calibration source is known to within 1.5 cm. Additionally, a tilt of the calibration deployment system relative to the IV is known to be small (< 5 degrees), but is determined much more precisely using the laserball source described in Sec. VIII.

The PMT cables are plugged into custom high voltage boards (HVSSs) that extract the PMT pulses and pass them to a set of CAEN V1730 digitizers [27]. The HVSSs are also responsible for producing an approximately 16 ns wide, 20 mV square pulse for each PMT pulse that crosses a programmable threshold. These square pulses are used for triggering the detector. The HVSSs each perform an analog sum of the square pulses across sixteen channels, and pass the sum to a central analog summing board (CASB). The CASB performs a global analog sum of the trigger signals and compares the sum against a configurable global threshold. The CASB provides five thresholds across two gain paths with different dynamic ranges. If any of the thresholds are crossed, a trigger is emitted and sent to the central trigger board (CTB). The CTB hosts an FPGA, which is used to apply programmable trigger logic to the trigger inputs. For example, the CTB allows one to trigger on the coincidence between any of the five thresholds within a certain time-window. If the trigger condition on the CTB is met, a final “global trigger” is emitted and sent to the CAEN boards to inform them to read out their data to disk. In addition to the CAEN data, trigger information from the CTB is also written out. The CTB provides a clock that is fanned out to the CAEN boards in order to keep the separate electronic components synchronized.

All 241 PMT signals are digitized using CAEN V1730 waveform digitizers at a sampling rate of 500 MS/s. The data is readout over optical link and stored in *hdf5* files for processing. A description of the complete downstream processing chain can be found in Sec. IV.

III. DETECTOR MODELING

The Eos detector is modeled using the open-source, publicly available RAT-PAC-2 simulation and analysis package [28]. RAT-PAC-2 contains Geant4-based [29] detector geometry descriptions, custom optical photon generation and propagation, individual 3D PMT modeling, a flexible database that holds run-dependent parameters, and customizable data-acquisition simulations. In Eos we simulate a detailed detector geometry, as visualized in Fig. 3, that includes all of the PMTs, dichroicons, and material within the detector. The calibration sources are also modeled in similar detail. The physics interactions are handled by Geant4 and are initiated according to the details of each Eos calibration source. Geant4 version v11.1.2 is used for the simulation. Standard particle interactions are simulated using the “shielding” reference physics list in Geant4




FIG. 3. The Eos detector geometry model as implemented in the RAT-PAC-2 software. The various geometry features include the three different types of PMTs, the dichroicons (pink), the inner IV (cyan), the acrylic neck, and the steel support structures.

[30], which is optimized for low energy interactions and neutron transport. Optical photons are generated and propagated through the detector using custom optical physics modules, which are based on standard Geant4 processes but include improvements to the handling of scintillation photon generation, absorption, and Rayleigh scattering.

The PMT response model is based on ex-situ and in-situ measurements of the quantum efficiency, transit time distribution, pulse shape, and pulse charge. The cable delays and set of offline PMTs are monitored on a regular basis and included into the RAT-PAC-2 database. Further details about how the PMT parameters are extracted are provided in Sec. VIII.

A simplified simulation of the Eos trigger performs an analog sum of the square pulses created for each PMT hit. This trigger decides when an event has crossed threshold and flags RAT-PAC-2 to write the data to disk. Typically, the threshold is selected to be sufficiently low such that we trigger on more events in the simulation than in the data. This allows for the flexibility to study below-threshold events. An analysis-level threshold is applied to both the data and simulation, as described further in Sec. VIII.

For triggered events, each of the above-threshold pulses are passed to a model of the CAEN V1730 waveform digitizer with appropriate settings (fetched from a run-level database). This model digitizes the pulses and passes the set of digitized waveforms to a waveform analysis tool.

This tool is identical to the one used to process the data, described in Sec. IV. The outputs of the waveform analysis are the extracted hit times and charges, which are passed to the reconstruction algorithms described below. This allows for a straightforward, direct comparison between the simulation and data.

IV. DATA PROCESSING

A complete processing chain, written as part of RAT-PAC-2, reads the `hdf5` files produced by the data acquisition system (Sec. II) and extracts PMT times and charges that can then be used as inputs to reconstruction and analysis tools. The data processing tool generates detector data in the same format as the simulated events, allowing them to be analyzed in identical ways.

In order to extract the pulse times and charges from the waveforms, above-threshold waveforms are identified by applying a negative 5-mV threshold. If a channel has a downward-going pulse that crosses below threshold, the time associated with the peak is identified and the PMT gets counted toward the global “hit” sum (N_{hit}). In this manuscript, we refer to the total number of waveforms with a peak that crosses threshold as the N_{hit} .

After the initial N_{hit} calculation, a cross-talk cut is applied to remove unwanted, bipolar signals from inductive pickup. An example of such a signal is shown in Fig. 4 (left). This cut is particularly necessary for high-energy cosmic events, where large currents in a single channel can induce threshold crossing waveforms on many neighboring channels. The cross-talk cut is defined by the ratio of the negative to positive amplitude of each pulse. For induced signals this ratio is close to unity, while for pulses from photoelectrons the ratio is significantly larger. Single photoelectron laserball data (Sec. VA) was used to determine the optimal cut value, which provided a clear separation between photoelectron signals and cross-talk.

Waveforms that do not pass this cut are removed from the N_{hit} of the event.

A lognormal function of the following form is fit in a range of $[-10, +15]$ ns around the peak of the above-threshold pulse:

$$f(t) = B + \frac{Q}{(t - t_0)\sqrt{2\pi}\sigma} e^{-\frac{1}{2}(\log(\frac{t-t_0}{m})/\sigma)^2}, \quad (1)$$

where B is the baseline of the waveform, Q is the charge contained in the pulse, t_0 is the arrival time of the pulse, σ is the shape parameter, and m is a scale parameter that sets the median of the distribution. In the fit, the values of m , B , t_0 , and Q are floated and the value of σ is fixed to 0.15. Figure 4 (right) shows an example waveform from an R14688-100 PMT, with the associated lognormal fit. The fit captures the shape of the PMT pulses in the data and the χ^2/NDF values are typically around one. The PMT time is calculated from the fit using

$$t_{\text{PMT}} = t_0 + m. \quad (2)$$

Other methods to extract the pulse timing have been explored, including applying a constant-fraction discriminator, fitting a Gaussian function, and applying intersample sinc interpolation (see the methods described in Ref. [31]). The lognormal fit achieves better or equivalent timing when studying calibration data, and is the method used to extract the PMT times in this manuscript.

In the data collected for this paper, the PMTs primarily detect single photoelectrons (SPEs). Therefore, the lognormal fits extract only a single hit-time for each waveform. Other methods using machine learning (ML) [32] and other sophisticated approaches [33–35] are currently under development. These tools will be appropriate for assigning hit-times and counting photoelectrons (PEs) for waveforms with multiple photoelectrons (MPEs).




FIG. 4. Left: an example digitized waveform from pickup that passes the 5-mV hit threshold, collected during a cosmic run. This hit is removed by the cross-talk cut described in Sec. IV. Right: an example digitized waveform for an R14688-100 PMT from Eos laserball data. The red lognormal fit is performed according to Eq. (1) and is used to extract the arrival time of the pulse.

V. DATA COLLECTION

The data collected during the Eos water phase consist of deployed optical and radioactive sources. These sources are lowered into the detector via a calibration control system capable of deploying sources at fixed locations along the central vertical axis (z). In addition to deployed sources, Eos collected cosmic data to tag Michel electrons, which provides a well-understood, uniformly distributed source of events. Eos has also deployed several additional sources not described in this manuscript, including an internal acrylic Cherenkov source and an external γ source. The sources each have associated detector settings that get loaded at the start of each data collection “run.” These settings include the trigger threshold and trigger scheme (e.g., if a coincidence trigger is required). Table I summarizes the run plan for each of the calibration sources described in this manuscript. The total number of events collected and position of the sources were determined on a source-by-source basis based on calibration and reconstruction testing needs. Typically, several runs per calibration source position are collected, and repeated runs are used to constrain systematics. Significantly more data (mostly with the laserball) than presented in this manuscript has been collected with Eos, but it has been primarily used for ensuring the stability of the detector performance with time. The activity of all radioactive sources is low enough such that we expect negligible pile-up (more than one decay in a 280 ns trigger window).

A. Laserball source

The laserball source provides isotropic light at a fixed wavelength and tunable intensity. This is achieved by pulsing one of four lasers (374 nm, 408 nm, 442 nm, and 515 nm) into a single-mode fiber that is connected directly into a 9.5 mm diameter diffusing ball. The laserball is made of Teflon which scatters the light sufficiently to create an isotropic output. While none of the measurements in this manuscript rely on the laserball being perfectly isotropic, the deviations from a perfectly isotropic source are tested on the benchtop and show less than 5% variations as a function of output angle.

The intensity of the laser is tuned by measuring the coincidence rate (between the laser pulse and a PMT pulse)

for any given PMT and requiring that it is below 5%. This ensures that the light detected at the PMTs is highly SPE. The pulsed trigger for the laser is used to trigger the digitizers, ensuring that the PMT light arrives back to the digitizers at a fixed time after the trigger signal. The laserball is used for most of the detector calibrations described in Sec. VIII.

B. Thorium source

The thorium source consists of ten cylindrical tungsten rods, each 17.5 cm long with a radius of 2.4 mm, arranged in a ring inside a cylindrical container of black Delrin. The tungsten rods contain 4% ^{232}Th by mass. The ^{232}Th decay chain produces several β , α , and γ emitters, the former two of which are absorbed in the opaque Delrin. The γ -rays exit the source and Compton scatter in the detector, producing Cherenkov light. The primary γ -rays come from ^{208}Tl (2.6 MeV). Lower energy γ 's are produced by ^{228}Ac , ^{212}Bi , and other isotopes, altogether resulting in more than 10 times fewer observable interactions in the detector than the ^{208}Tl γ -ray, based on RAT-PAC-2 simulations. The thorium source provides a sample of nearly monoenergetic γ 's with a known original position distribution (from within the source), allowing for the study of position and direction reconstruction performance.

To trigger on the thorium source data, a trigger threshold on the CASB is set. This threshold is selected to study the spectrum of events emitted by the thorium source, without being so low that the dark-noise and other PMT related backgrounds start to dominate. A threshold of 100 mV was selected based on studies of the trigger efficiency as a function of N_{hit} , described in Sec. VIII E. This threshold roughly corresponds to five in-time (within the 16 ns trigger pulse width) N_{hit} . To sustain the trigger rate at this threshold the CTB prescales the data by a factor of 32 (i.e., only 1 in 32 triggered events are collected), which yields a trigger rate around 500 Hz while the thorium source is deployed.

C. Directional β sources

The directional β sources, shown schematically for one design in Fig. 5 (left), consists of β -emitting isotopes surrounded by metal and delrin shielding with boreholes

TABLE I. Details of calibration source deployments in Eos. The z range and step size indicate where data were collected along the central axis of the detector. The number of events collected is per calibration source position and was selected based on the trigger scheme and the relevant background rates. The laserball scan was collected with four different wavelengths, described in Sec. VA. The directional source data was collected with two different radioactive sources, ^{90}Sr and ^{106}Ru , as described in Sec. VC.

Source	z range (cm)	z step (cm)	Events per run (10^5)	Trigger scheme
Laserball	[−60, 60]	10	2	Pulsed
Thorium	[−50, 50]	10	5	Single threshold
AmBe	[−30, 30]	30	10	Coincidence
Directional	[−50, 50]	Various	0.5	Self-triggered




FIG. 5. Left: a schematic of the 30 mm directional source that shows the internal components necessary for shielding and self-triggering. Right: design of the 30 mm directional source in RAT-PAC-2.

that allow β 's to exit the sources as a collimated stream in a fixed direction. Two different isotopes are used, ^{90}Sr and ^{106}Ru . The two isotopes emit β 's with Q -values of 2.28 MeV and 3.5 MeV respectively, with ^{106}Ru emitting deexcitation γ s for 20% of the decays. The sources can be rotated to desired zenith angle (θ) and azimuthal angle (ϕ). The RAT-PAC-2 model of the source is shown in Fig. 5 (right).

A self-triggering system, consisting of a scintillating fiber ribbon and two silicon photomultipliers (SiPMs), tags exiting β 's with an approximately 50% efficiency. Electronics inside the source discriminate the SiPM pulses and emit trigger pulses that are plugged directly into the CTB. The CTB is configured to emit global triggers when the directional source self-trigger occurs in coincidence with a CASB-generated trigger. A 128 ns window is used for this coincidence, which was tuned to provide a pure sample of tagged directional source events. The CASB threshold for the directional source data is typically set to 50 mV. The rate of exiting β 's out of the source into the detector is on the order of a few Hz and the coincidence between the self-trigger and the detector trigger provides a clean dataset with less than 1% background.

The source is constructed in three different sizes to study the effect of the shadowing caused by the source body in future scintillator phases. The results described in this paper use both 30 mm sources and the 20 mm ^{106}Ru source. These directional sources provide a sample of β events with a known position, energy distribution, and direction, making them useful for studying reconstruction performance, particularly direction reconstruction.

D. AmBe source

The $^{241}\text{Am}^9\text{Be}$ (referred to as AmBe) source produces around 900 Hz of neutrons, more than half of which are

accompanied by a 4.4-MeV γ . An α from the decay of ^{241}Am is absorbed by ^9Be to produce $^{12}\text{C}^*$ and a neutron, where the $^{12}\text{C}^*$ deexcites with the 4.4-MeV γ . The neutron will capture on a H atom in the water and produce a 2.2-MeV γ with a mean capture time of around 210 μs . The delayed coincidence of these two γ 's allows for the pure selection of AmBe signals.

^{241}Am also produces a large number of X-rays, most abundantly at 59 keV. To absorb most of these, the AmBe source is enclosed in 1.5 mm of lead. Furthermore, the neutrons produced can have energies up to around 10 MeV, which can induce scintillation light (when the target is a scintillator) via elastic scatters with protons. To prevent contamination of the prompt 4.4-MeV γ signal by these scatters, the AmBe source is further enclosed in a large cylinder of opaque black acrylic with a height of 13 cm and a diameter of 14 cm. Thus, the AmBe source has the potential to provide two monoenergetic γ 's with a known origin and an understood distribution, allowing the study of position, direction, and energy reconstruction performance.

The trigger scheme for the AmBe relies on the coincidence between the multiple gain paths from the CASB, as performed by the CTB. A threshold of 800 mV is configured for the lowest gain path, which is used primarily to tag muons passing through the detector. Any event that causes this gain path on the CASB to trigger is excluded and no global trigger is emitted. Additionally, all triggers are locked out for approximately 1 μs after a high threshold trigger on the low gain path, to reduce contamination produced by ringing, after-pulsing, and short-lived followers after a cosmic muon interaction.

A threshold of 120 mV is set on the medium gain path (used for the prompt events) and a threshold of 70 mV is set on the high gain path (used for the delayed events). The logic on the CTB waits for a trigger on the medium gain path (with no accompanying trigger on the low gain path)




FIG. 6. Left: the N_{hit} of a central thorium source run compared to the associated background run, between 10 and 80 N_{hit} . The data for the background run is scaled by the relative live time. Right: the reconstructed ρ^2/ρ_{IV}^2 distribution is from the SeedNDestroy vertex fitter (see Sec. VI A), for central thorium and background data, selecting events in an N_{hit} range between 15 and 40.

and then masks in the high gain path for 400 μs . More precisely, the high gain (low threshold) path is enabled in a window of [1, 401] μs after a medium gain trigger. Excluding the first 1 μs rejects any reflections or other noise on the trigger without significantly impacting the efficiency of the selection. This trigger scheme is specifically designed to select the prompt γ -rays from the AmBe (4.4 MeV signal) succeeded by the neutron capture on a H atom (2.2 MeV signal), while attempting to reduce backgrounds from cosmic muon followers. The selection of the thresholds for the prompt and delayed event are configured as low as possible such that triggering without a prescale is possible.

E. Cosmic ray muons

Due to minimal overburden, Eos is bombarded by cosmic ray muons at a rate of several kHz. These muons and the spallation products created by them can cause significant backgrounds across a wide range of energies. However, they also provide an opportunity to act as an additional calibration source. In particular, both Michel electrons and neutrons are easy to tag by looking for high energy events followed closely in time by a second event.

To trigger on cosmic muons and their followers, a primary threshold of 800 mV, equivalent to roughly 40–50 N_{hit} , is used to tag high energy events. Similar to the trigger scheme for the AmBe, a low energy follower window is opened for 400 μs with a threshold of 140 mV, excluding the first 1 μs after the primary trigger. This high energy follower trigger provides a sample of easily tagged Michel electrons.

F. Backgrounds

Other than the directional source, the radioactive sources described in this paper are not tagged sources (e.g., by a sensor inside the source). Therefore, Eos relies on the

relatively high decay rates of the calibration sources to overcome the background rates, for example, from radioactive decay products of uranium and thorium in the water or PMT glass or from spallation products from muons. To account for the effect of the backgrounds on the measured quantities, dedicated runs are collected with no source in the detector and with the detector settings configured identically to the associated calibration source run. For example, background runs are collected directly before and after the thorium source data collection campaign with detector settings identical to the thorium source runs. The N_{hit} and ρ^2/ρ_{IV}^2 (where $\rho = \sqrt{x^2 + y^2}$ and $\rho_{IV} = \sqrt{91.2^2 + 91.2^2}$ uses the dimensions of the IV, in centimeters) of a central thorium run ($z = 0$) are compared to the associated background run in Fig. 6. The background distributions are scaled using the relative data collection live times. In the radioactive source analyses, the background data is statistically subtracted from the calibration source data. In several of the analyses, a fiducial volume cut is used to reject events at high radii, which are predominantly background.

VI. EVENT RECONSTRUCTION

As part of the primary detector goals, Eos is being used to develop and test both traditional and novel reconstruction methods utilizing both the Cherenkov and scintillation light in a ton-scale hybrid detector. Toward this goal, the reconstruction methods have been developed to run in a pure Cherenkov detector and were characterized during the water phase. In this section we present the details of each of the algorithms. In Sec. XI we detail the reconstruction performance for each of the various reconstruction algorithms, and show data and simulation comparisons.

Reconstructing events in Eos poses unique opportunities and challenges that will help develop methods for future detectors, such as Theia. Considering only Cherenkov light

for this manuscript, the reconstruction challenges include an asymmetric detector orientation (e.g., different PMTs at the top and bottom), large calibration source sizes (relative to the size of the detector), and the relatively large sensor to detector size ratio.

The reconstruction methods described in this section all share several features. First, all take as input the PMT timing information produced by the lognormal fitting procedure performed over the waveforms, as described in Sec. IV. All methods use the same set of hit PMTs, including the removal of cross-talk channels, and assume all hits are SPEs. Because the water data is dominated by SPE hits, no PMT charge information is currently used. The likelihood algorithms and the ML-based algorithm are flexible to enable the use of PMT charge information in the future. Currently, only the ML-based algorithm, described in Sec. VIC, simultaneously fits for the vertex and direction, and upgrades to the likelihood-based approaches could enable them to fit for both event position and direction at the same time.

The likelihood-based fitting methods use simulated events to produce probability density functions (PDFs). In all cases, a set of PDFs is generated using simulated electrons, distributed uniformly in the detector, with isotropically selected directions, and ranging in energies from 0.5 to 15.0 MeV. A total of twenty million events are simulated for PDF construction and for training the ML algorithm described in Sec. VIC. This method for PDF generation was selected because MeV-scale neutrino interactions are typically detected in Cherenkov detectors through neutrino-electron scattering, and thus are reconstructed under the hypothesis that the event is an electron. By simulating uniformly in position and direction and across a wide energy range, we avoid biasing the fit methods. Other methods for generating the PDFs were investigated, such as using simulations of the calibration sources themselves, and have only a very small impact on the reconstruction results.

A. Vertex reconstruction

There are several algorithms, described below, that run as part of the Eos analysis chain to produce reconstructed event vertices. In all cases, the event position is defined by three Cartesian coordinates, with z as the vertical axis of the detector, an event time [t_{recon} in Eq. (3)], and the origin (0, 0, 0) of the detector at the center of the IV. An important quantity that is calculated and used by reconstruction methods is the PMT time-residual. These residuals for a hit PMT are calculated as

$$t_{\text{res}} = t_{\text{PMT}} - t_{\text{tof}} - t_{\text{recon}}, \quad (3)$$

where t_{PMT} is extracted using the lognormal fitting procedure and is defined in Eq. (2), t_{tof} is the photon time-of-flight from the reconstructed position to the PMT,

and t_{recon} is the reconstructed time of the event. In general, t_{recon} captures a DAQ- and trigger-dependent arbitrary offset of the time-residuals from zero. To calculate the time-of-flight, the speed of 400-nm photons in water is set to 222.69 mm/ns (using $n = 1.344$) [36] and straight paths are assumed. The value at 400 nm is used because it roughly corresponds to where the PMT quantum efficiency distribution peaks, and selecting other reasonable fixed values of n has negligible impacts on the results.

In certain scenarios (e.g., for some calibration methods), the known calibration source position, rather than the reconstructed position, is used to generate the time-residuals, in which case the value of t_{recon} is set to zero and the time-residuals are indicated as $t_{\text{res}}^{\text{source}}$.

1. Quad fitter

The “quad fitter” runs on events with four or more PMT hits. The method selects four PMTs from the set of hit PMTs and calculates the event vertex based on the hit-times of those four PMTs, assuming the light traveled in a straight line from a single interaction point in the detector. This method is then repeated for a new set of four PMTs to produce a second event vertex position. A “cloud” of estimated vertex points is generated and the median of that cloud is selected as the event vertex. The quad fitter has the advantages of being fast and very robust, and is thus used to seed the position of some of the slower, more precise fitting algorithms. This method has been used by large Cherenkov experiments, such as SNO [37].

2. Likelihood-based approaches

Eos utilizes two independently developed likelihood-based vertex fitters called SeedNDDestroy and Mimir. These frameworks perform likelihood-based event reconstruction strategies in water Cherenkov and scintillation detectors. The methods both reconstruct the position of the event by minimizing a PDF in time residual. Formally, the vertex $\theta = (\vec{x}, t)$ is estimated by maximizing the following likelihood:

$$\mathcal{L}(\theta|\{t_{\text{res}}^i\}) = \prod_i \mathcal{P}(t_{\text{res}}^i), \quad (4)$$

where t_{res}^i is the time residual of the i th hit of the event. In this manuscript, SeedNDDestroy uses a single PDF in time residual and Mimir uses distinct PDFs for each PMT type (R14688-100, R11780, and R7081). Both likelihood fitters are seeded with the result from the quad fitter. The optimization used by SeedNDDestroy is the NLOpt [38] implementation of the Subplex algorithm, Sbplx [39]. Mimir uses the ROOT optimizer MINUIT2 [40]. Because Mimir and SeedNDDestroy perform similar calculations they are expected to yield similar results in Eos. Having two similar algorithms is helpful for providing robustness checks and identifying bugs and outliers in individual algorithms.

Additionally, Mimir internally continues to fit for the event direction (Sec. VIB) using the vertex fitting result as the hypothesis.

B. Direction reconstruction

As with the algorithms that are used to fit event vertex, there are several algorithms, described below, that run as part of the Eos analysis chain to produce a reconstructed event direction. The performance of these algorithms is primarily characterized on the directional sources described in Sec. VC. For these sources, the reconstruction of event direction is studied by calculating:

$$\alpha = \hat{d}_{\text{source}} \cdot \hat{d}_{\text{recon}}, \quad (5)$$

where $\hat{d}_{\text{source}}$ is the known direction in which the source is pointing and $\hat{d}_{\text{recon}}$ is the reconstructed direction. This metric is appropriate for the directional source, where the average emission angle of the electrons from the source is known. We expect to find a peak in α at a value of one. This metric does not provide an absolute measure of reconstruction resolution in Eos. Instead, a comparison of data to simulation allows us to validate the model, which can then be used to evaluate direction resolution for a defined event type.

The distribution of α over many reconstructed events tends to be exponential and can be quantified more easily by fitting an empirical function to the α distributions. The function form, f , is an exponential convolved with a Gaussian where the exponential time-constant and the Gaussian width and mean are floated in the fit. After performing this fit, the 1σ (68%) directional resolution ($\alpha_{1\sigma}$) is calculated by finding the value when $\int_{\alpha_{1\sigma}}^1 f d\alpha = 0.68$. Conceptually $\alpha_{1\sigma}$ quantifies the fraction of events that reconstruct with directions similar to the true direction. The fit also provides the uncertainties on $\alpha_{1\sigma}$.

For each algorithm the direction reconstruction is calculated using only hits in a prompt time-residual window (between -3 and 5 ns), where the direct Cherenkov photons are expected to arrive. The hits outside this window are likely from PMT dark current or scattered/reflected photons, the inclusion of which would smear out the direction resolution.

1. Average direction

The average direction fitter takes the average direction of the vectors that point from an input event position to the hit PMT positions. The average direction algorithm has the advantages of being fast and robust, and thus can be used to seed the direction of the likelihood-based fitting algorithms. The event position input for the average direction fitter is taken from the SeedNDDestroy result.

2. AERO

The likelihood-based fitter angular event reconstruction via likelihood optimization (AERO) utilizes HEALPix [41], an algorithm designed to discretize a spherical surface into equal-area pixels, allowing uniform sampling of directions in three-dimensional space. To create the PDFs for the likelihood-based optimization, events are binned into appropriate HEALPix pixels based on the true initial direction of the electron. For each of the 192 HEALPix pixels a 2D PDF of time-residuals vs PMT ID is generated. Conceptually, these PDFs represent the PMT timing and occupancy information as a function of the event direction (binned into 192 directions on a sphere).

When evaluating the likelihood for a hypothesized direction, AERO averages the PDFs of the 20 nearest pixels, weighted by their angular distances to the evaluation direction, to generate a new, interpolated PDF. This interpolation method allows the reconstruction of events across the full sphere, rather than only at 192 fixed directions. The likelihood of an event in the evaluation direction producing PMT and time-residual patterns seen in the PDF is calculated. The direction that minimizes the negative log-likelihood is selected as the best-fit hypothesis. AERO uses the ROOT optimizer MINUIT2 and an input event position from the quad fitter result.

3. Mimir

Following the general framework and vertex fitting routine described in Sec. VIA 2, a direction fitting strategy is performed. In this step, the direction of the particle $\theta_{\text{dir}} = \hat{d}_{\text{Mimir}}$ is estimated by maximizing the likelihood:

$$\mathcal{L}(\{\hat{d}^i\}|\theta_{\text{dir}}) = \prod_i \mathcal{P}_{\text{type}}(\hat{d}_{\text{Mimir}} \cdot \hat{d}^i), \quad (6)$$

where $\hat{d}^i$ is the unit vector pointing from the event vertex to the position of the i th PMT hit in the event.

C. Deep-learning-based reconstruction

Highly parallelizable inference technique for monitoring antineutrinos (HitMan) is a simulation-based inference tool that uses deep learning techniques to generate likelihood contours for particle kinematics θ given photosensor data $\mathbf{x}$ [42]. It provides a surrogate model for the detector likelihood using neural ratio estimation (NRE).

1. Neural ratio estimation

Unlike regression models that typically provide point estimates, the NRE learns the entire likelihood function. This is accomplished by training a classifier to distinguish correlated simulation pairs $\{\mathbf{x}, \theta\}$ from uncorrelated pairs, approximating the likelihood-to-evidence ratio [43]. To efficiently handle sparse, variable-sized PMT data, we

decompose the likelihood using an extended maximum likelihood formulation [42]:

$$\mathcal{L}(\boldsymbol{\theta}|\mathbf{x}) = \left[\prod_{i=1}^{N_{\text{tot}}} q(x_i|\boldsymbol{\theta}) \right] \mathcal{P}(N_{\text{tot}}|\Lambda_{\text{tot}}(\boldsymbol{\theta})). \quad (7)$$

Here, $\mathcal{L}(\boldsymbol{\theta}|\mathbf{x})$ is the likelihood function with real data $\mathbf{x}$ under some hypothesis $\boldsymbol{\theta}$, N_{tot} is the total observed photons in the entire detector with expected value $\Lambda_{\text{tot}}(\boldsymbol{\theta})$, and $q(x_i|\boldsymbol{\theta})$ is the probability density for a single photon measurement i with features x_i (detection position and time).

Two NRE networks are trained to approximate $q(x_i|\boldsymbol{\theta})$ and $\mathcal{P}(N_{\text{tot}}|\Lambda_{\text{tot}}(\boldsymbol{\theta}))$:

- (i) HitNet: Approximates the per-photon term $q(x_i|\boldsymbol{\theta})$, processing individual photon detection features relative to $\boldsymbol{\theta}$.
- (ii) ChargeNet: Approximates the Poisson-like term $\mathcal{P}(N_{\text{tot}}|\Lambda_{\text{tot}}(\boldsymbol{\theta}))$, modeling the total number of detected photons (does not use charge in this manuscript, but is capable of doing so in the future).

The total log-likelihood is the sum of the ChargeNet output and the HitNet outputs summed over all photons.

2. Implementation and parameter estimation

For the water phase analysis, the hypothesis $\boldsymbol{\theta}$ includes seven parameters: initial vertex (x , y , z), time (t_{recon}), direction (ϕ , ϑ), and energy (E). The networks are comprised of dense layers with Mish activations [44], trained on 20 million simulated electron events (0.5–15.0 MeV) described in Sec. VI.

Network inputs are feature-engineered; for example, HitNet uses the time residual between the photon detection and hypothesis so the network obeys event time translation symmetries. Currently, HitNet and ChargeNet rely on the single photoelectron assumption dictated by the current waveform analysis, meaning each hit PMT corresponds to a single photon detection ($N_{\text{tot}} = N_{\text{hit}}$). This approximation is sufficient for energy depositions up to a few MeV in water. Once MPE extraction is implemented (Sec. IV) these results will be naturally included in the current framework.

For parameter estimation, we find the maximum likelihood estimate (MLE) $\hat{\boldsymbol{\theta}}$ by minimizing the negative

log-likelihood generated by HitNet and ChargeNet. A seed from the Quad fitter (Sec. VIA) is refined using the gradient-free Sbplx algorithm.

VII. EVENT SELECTION

Several techniques are used in Eos to reject ambient and instrumental backgrounds. For example, fiducial volume cuts, to select events near the center of the detector, avoiding the regions of the detector that are less well-understood and contain higher rates of backgrounds. The procedure to statistically subtract background distributions, described in Sec. VF is also applied for several sources. This section defines additional event selection criteria, which are specific to each source. The cuts are described in the following sections and summarized for each source in Table II.

A. Prompt N_{hit} selection

It is useful in Eos to define a prompt N_{hit} ($N_{\text{hit}}^{\text{p}}$). The prompt N_{hit} criterion selects hits in a narrow time window around the peak of the time-residual distribution, rather than counting the total number of hits over the full 220 ns trigger window. Using $N_{\text{hit}}^{\text{p}}$ for calibrations greatly reduces the impact of mismodelings of late light from the PMT transit response or from reflections in the detector. In this manuscript, $N_{\text{hit}}^{\text{p}}$ is determined by selecting hits $-2 \text{ ns} < t_{\text{res}}^{\text{peak}} < 4 \text{ ns}$, where $t_{\text{res}}^{\text{peak}}$ is the peak of the time-residual distribution (which is typically centered at 0 ns). Different $N_{\text{hit}}^{\text{p}}$ cuts are applied for each dataset, depending on the energy region of interest. These cuts are necessary for several of the radioactive sources in order to avoid the trigger inefficiencies described later in Sec. VIIE, which are not modeled in detail.

B. Directional source selection

As described in Sec. VC, the directional sources use a self-triggering system consisting of a SiPM and scintillating fiber ribbon at the tip of the source. The time difference between the self-triggering time and the average PMT hit time for each event is calculated. A Gaussian fit to a histogram of this time-difference, across the full datasets, yields an average expected time-difference and events

TABLE II. A summary of the various analysis cuts for each of the calibration sources. Note that the directional sources also utilize the coincidence cut described in Sec. VIIB.

Source	$N_{\text{hit}}^{\text{p}}$	$ z $ (cm)	ρ (cm)	Δr (m)	Δt (μs)	Background subtract
Laserball						
Thorium	> 10	< 70	< 70			✓
Directional	> 10	< 70	< 70			
AmBe prompt	[23, 39]	< 70	< 70			✓
AmBe delayed	[9, 21]	< 70	< 70	< 1	< 400	✓
Michel e^-	[51, 99]	< 70	< 70		> 1 and < 5	




FIG. 7. Left: the N_{hit} distribution for the 20 mm ^{106}Ru directional source data before and after applying the directional source selection cuts described in Sec. VIB. The peak around 200 N_{hit} is from accidental coincidences from cosmic muons, which are mostly removed through the directional source event selection. The uncertainties are statistical only. Right: event display of a tagged β event from the ^{90}Sr directional beta source, located at the center of the detector and pointed sideways, toward the barrel PMTs. The color scale indicates the charge of prompt hits, capped at 2 pC.

within $\pm 2\sigma$ are accepted. This cut removes the majority of the accidental coincidences between the detector trigger and the self-trigger from the directional sources. Figure 7 (left) shows the N_{hit} distribution for the 20 mm ^{106}Ru directional source before and after applying the cut on the time difference. Figure 7 (right) shows the Eos event display for the sideways pointed directional source and a clear Cherenkov ring is visible.

C. AmBe source selection

A special coincidence tag is developed for the AmBe source, described in Sec. VD. The prompt-event candidates, from the 4.4 MeV $^{12}\text{C}^*$ deexcitation γ -rays, are selected using an $N_{\text{hit}}^{\text{p}}$ range of 22 to 40 and with a reconstructed radial position of $r < 75$ cm (using SeedNDestroy). For each prompt event candidate, delayed events are identified in a subsequent 400 μs window with $8 < N_{\text{hit}}^{\text{p}} < 22$, $r < 75$ cm, and radial distance from the prompt event of $\Delta r < 1$ m. The cut values were informed using the simulation to maintain a high efficiency for selecting coincidences, while removing the vast majority of accidental backgrounds. The low value of the cut on $N_{\text{hit}}^{\text{p}}$ for the delayed event is acceptable because of the trigger scheme described in VD, which considerably lowers the trigger threshold on delayed events.

If either a promptlike event or a high- N_{hit} (above 70 N_{hit}) muonlike event is identified in the delayed event window, that prompt-event candidate is rejected. In situations with multiple delayedlike events in the follower window, all delayed events are accepted. The high rate of low-energy

events in Eos, due to its location on the surface and the relatively high levels of radioactivity, make it unfeasible to apply a multiplicity cut to remove events with multiple delayedlike events. In a deep-underground, clean detector the event selection would be simplified due to significantly lower backgrounds.

Figure 8 (left) shows the time difference (Δt) between the selected prompt and delayed events for a central AmBe run. Good agreement between data and simulation is observed and a fitted neutron capture time yields $\tau = 202.3 \pm 6.4 \mu\text{s}$, consistent with expectations. The 400 μs cutoff is caused by the size of the coincidence window implemented in the triggering logic (for details, see Sec. VD). This indicates that the coincidence between the prompt 4.4 MeV γ and delayed neutron capture, producing a 2.2 MeV γ can be selected. Figure 8 (right) shows the Eos event display for a tagged prompt event from the centrally deployed AmBe source.

D. Michel electron selection

Michel electrons can be efficiently tagged in Eos using the high-energy follower trigger described in Sec. VE. Events above 100 $N_{\text{hit}}^{\text{p}}$ are classified as muon candidates and Michel electron followers are searched for in a subsequent 10 μs window. The Michel electron $N_{\text{hit}}^{\text{p}}$ window is selected between 50 and 100 based on MC studies. A multiplicity cut ensures exactly one candidate Michel electron occurs after the high-energy trigger in a 10 μs window. The Δt after the prompt muon is required to be larger than 1 μs , which removes reflections and other noise




FIG. 8. Left: the time-difference between the prompt and delayed events for a central AmBe run, shown for both data and MC. An exponential fit to the data yields the expected neutron capture time. The uncertainties are statistical only. Right: event display for a tagged prompt event from the central AmBe run. The color scale indicates the charge (in pC) of prompt hits and the dichroicon PMTs at the bottom are indicated by the dashed circles.

in the trigger system, prevalent primarily for high-energy events. Additionally the requirement that the $\Delta t < 5 \mu\text{s}$ selects a higher purity of Michel electrons.

Figure 9 (left) shows the Δt between the cosmic muons and the Michel electron followers in a timing window of $10 \mu\text{s}$. The timing spectrum is fitted to an exponential decay function plus a constant background. This yields a muon lifetime of $\tau = 2.05 \pm 0.08 \mu\text{s}$ with $\chi^2/\text{d.o.f.} = 24/19$. The expected muon lifetime is calculated in Appendix B to be $2.04 \pm 0.01 \mu\text{s}$, which is in

agreement with the fitted result. Figure 9 (right) shows the event display of a Michel electron event detected with the high-energy follower trigger, in which a clear Cherenkov ring is identified.

E. Summary

Table II summarizes the analysis selections for each of the calibration sources. Identical selection criteria were applied on both the simulated and measured datasets.




FIG. 9. Left: the time difference between selected cosmic muons and the Michel electron followers, which is fit with an exponential plus a constant background. The uncertainties are statistical only. Right: event display of a Michel electron candidate. The color scale indicates the charge (in pC) of prompt hits and the dichroicon PMTs at the bottom are indicated by the dashed circles.

VIII. DETECTOR CALIBRATION

The calibration of the detector response is critical for tuning the simulation to better understand the Eos reconstruction performance in water and to prepare for future scintillator phases. This section focuses on the calibration of the PMT responses using the laserball and thorium sources. The output of these calibration efforts are used as inputs to the RAT-PAC-2 detector model. In Secs. X and XI, comparisons between data and simulation for several important observables are presented after applying these calibrations.

A. Timing delays

The duration between the time when a photon hits the photocathode and the time that is extracted from the digitized waveform varies from PMT to PMT. This variation is caused by different PMT transit times, cable lengths and path lengths along the electronics boards. This variation is calibrated individually for each PMT with laserball data collected using the 515 nm laser. It is necessary to use the 515 nm laser as the PMTs behind the dichroicons only detect enough light to perform this calibration at long-wavelengths (above 450 nm).

The time for each PMT signal is extracted using the method described in Sec. IV, accounting for the known position of the laserball source to subtract the photon time-of-flight. This time-residual distribution, combined across all PMTs, is shown in Fig. 10 (left).

To extract the per-PMT timing delay, the time-residual is calculated for each individual PMT over all the events in the laserball run. A Gaussian is fitted ± 0.5 ns around the peak of each PMT's time-residual distribution. The mean of the fit is selected as the timing delay for that PMT. The validity of the timing delay calibration is tested by applying the calibration to runs varying the z -position of the laserball

between 60 cm and -60 cm in steps of 10 cm. The result of applying this correction is demonstrated by plotting the time-residuals for all PMTs and for several laserball position in Fig. 10 (right). After applying the delay calibration for a central laserball run, a Gaussian is fitted ± 0.5 ns around the peak of the time-residuals. The σ of the fit to the central data is 450 ps. This resolution includes broadening effects from the laser jitter and scattering inside the laserball.

These timing offsets are applied directly to correct the data, and the simulation does not include any channel to channel differences in the delays. The time-residuals include the per-PMT cable delay corrections for all analyses in this manuscript.

B. PMT pulse shape

The PMT pulse shape is modeled as a Gaussian function in simulation. Using a Gaussian function (as opposed to a lognormal) allows us to more easily calibrate for the variation in pulse shape, and the skewness of observed PMT pulses is small enough that a Gaussian pulse model has minimal impact on fitted values such as charge and time.

A spectrum of fitted pulse widths is extracted from laserball data by fitting Gaussian functions and extracting the standard deviation, σ , from individual pulses. This σ spectrum is used in simulation to choose the shape of the simulated pulse on a hit-by-hit basis. The spectrum is adjusted to account for the broadening that occurs due to the process of refitting a function on a sampled pulse. It is then shifted by a multiplicative scaling factor S for each PMT.

It should be emphasized that this method is different than that used to extract the PMT hit-times (fitting a Gaussian rather than a lognormal distribution to the PMT pulse); however, the Gaussian model was simpler to implement




FIG. 10. Left: the laserball time-residuals for all PMTs in Eos prior to applying the timing delay calibration. The multiple peaks are in part due to the different cable lengths for the different PMT types. Right: the laserball time-residuals for all PMTs in Eos after applying the timing delay calibration. The broader distribution at larger z -positions is due to the proximity to the R11780 PMTs at the top of the detector, which have significantly broader transit-time spreads than the R14688-100 PMTs.

and resulted in a negligible difference in the final pulse-shape calibration results (when comparing to the results from the lognormal pulse shapes).

C. PMT charge

The PMT gain is set to 1×10^7 based on Hamamatsu recommendations for optimal performance. The high voltage delivered to each PMT to achieve this gain is determined from bench-top measurements of each PMT [23]. The PMT to PMT variation in the gain is calibrated *in situ* using laserball data. Two different charge distributions are studied: the charge integrated in a window $[-10, 20]$ samples around the peak of the waveform and the fitted charge, Q , from Eq. (1). For the purpose of the water phase, we focus on the latter charge and show here the results for the fitted charge. For future phases where MPE contributions are more significant, the integrated charge is expected to be more useful, and is calibrated similarly. As previously described, N_{hit} are counted as pulses that cross a threshold of -5 mV, so it is important to accurately model the PMT charge and pulse heights.

In the simulation the initial value of Q is selected by drawing from a histogram that is generated from bench-top SPE data collected for each PMT in the Eos detector [23]. A multiplicative scaling factor is introduced that shifts the value of Q on a hit-by-hit basis. This effectively shifts the Q distribution higher or lower in average charge. The value of this factor is adjusted using central laserball data. A histogram of Q is fitted with a Gaussian around the peak for both the data and simulation. The value of S is selected to shift the fitted mean value in the simulation to match the data. The distribution of Q for all PMTs for a central 408 nm laserball run, compared between data and simulation, is shown in Fig. 11 (left). After applying this calibration, the PMT pulse-heights are in good agreement between data and simulation, shown in Fig. 11 (right).

D. Dark rate

The PMT dark-rates are calibrated by looking in a pretrigger window in the laserball data. This early window is guaranteed to be prior to the laser firing, so contains only hits from the PMT dark noise. The digitizer windows are intentionally configured such that at least 80 ns of the window is prior to the trigger. The laserball source is selected for the dark-rate calibration as the data is collected every week, allowing a regular calibration of the dark rates.

A 60 ns window is selected for each PMT type (slightly different cable delays require slight differences in the windows) such that the window contains only the dark noise. The total number of PMT hits in that window for each PMT over all trigger events is calculated and converted to a rate. The R11780, R14688-100, and R7081 PMTs have a measured average dark rate of 8 kHz, 14 kHz and 4 kHz, respectively. These fixed rates are input into the simulation, where the dark rate is added as a uniform, random background for each PMT. Laserball runs collected throughout the dataset show that the dark-rate was stable over time.

E. Trigger efficiency

This efficiency of the trigger system as a function of the number of PMT hits can be studied using the laserball. The laserball operates by using a fixed rate, pulsed trigger that gets propagated to the CAEN digitizers regardless of the number of PMTs that detected light. The analog global trigger sum is digitized on an unused CAEN channel, and the pulse height of the sum is measured in an offline analysis. A map between the number of PMTs that detected light and the trigger sum is generated. This map can be used to measure the efficiency of the trigger, given a specified trigger threshold, as a function of the N_{hit} . It is assumed in the analysis that 100% of the pulses that cross the trigger threshold would have triggered the detector.




FIG. 11. Left: the data to simulation comparison for a central 408 nm laserball run after applying the charge calibration, for all online PMTs. Right: the same for the peak of the waveform.




FIG. 12. The trigger efficiency as a function of the N_{hit} for an assumed global trigger threshold of 100 mV (typical for a thorium source run).

Figure 12 shows the trigger efficiency as a function of N_{hit} for a threshold of 100 mV. The efficiency is calculated by finding the peak of the digitized global trigger sum, which is compared to a 100 mV threshold. If it falls below the threshold, it is assumed that no global trigger would have been generated, and vice versa. This assumption is based on bench-top testing and validation of the CASB trigger performance. Notably, for a 100 mV threshold, typical for thorium source runs, the trigger efficiency reaches 100% above 10 N_{hit} .

F. PMT efficiency

The wavelength-dependent PMT quantum efficiencies for each PMT type are extracted directly from the Hamamatsu datasheets [22]. The collection efficiency—the efficiency for which a PE created at the photocathode reaches the last dynode and produces a signal—is not

measured ex-situ for any of the PMTs. A scale correction to the overall detection efficiency, applied for each PMT type, is intended to calibrate the unknown collection efficiencies. These corrections can also account for other mismodeling that may contribute to a disagreement in the total amount of light collected by the PMTs, between the data and the simulation.

In order to calibrate these efficiency scales, a central deployment of the thorium source is used. The total number of photons emitted by the laserball is not known and is thus not appropriate for this calibration; however, the thorium source produces predominantly 2.6-MeV γ -rays that Compton scatter in the water and create Cherenkov light with a well-understood intensity.

For each PMT group, the resulting $N_{\text{hit}}^{\text{p}}$ histogram from the simulation is linearly interpolated into finer bins, and a least-squares fit is performed with the data to determine the scale factor that brings the simulation in to closest agreement. The fitted scale factors for each PMT group are then input into the simulation and the process is repeated until sufficient convergence is achieved. Three efficiency correction factors are obtained: 0.78 for the 168 barrel PMTs (R11468), 0.90 for the 24 top PMTs (R11780), and 0.88 for the 49 bottom PMTs (all bottom PMTs are grouped together due to their low N_{hit} and the more complex detector structure at the bottom). Figure 13 shows the resulting agreement between data and MC for all PMTs (left) and for the barrel PMTs (right) after applying the efficiency correction calibration.

G. Deployment system tilt

As described in Sec. II, a small tilt of the calibration tower relative to the IV is possible, based on the detector construction, and can be measured using the laserball source deployed at varying z positions. By evaluating the difference in the time-of-flights for opposite sides of the detector, it is determined that the laserball is traveling in




FIG. 13. The $N_{\text{hit}}^{\text{p}}$ distributions for a central thorium run, after applying the PMT efficiency calibrations. The left figure shows the data and simulation comparison for all PMTs and the right figure shows the barrel PMTs only. The uncertainties are statistical only.

the $(+x, +y)$ direction as it moves up the detector. The size of the tilt is measured to have a zenith angle $\theta = 0.65 \pm 0.35$ degrees. The uncertainty is evaluated by estimating the size of the tilt for several different laserball deployment scans. Due to the large uncertainty and lack of precision as to the orientation of the tilt in the azimuthal angle ϕ , this measurement is later used to estimate a systematic on the source position, and all simulations are performed assuming zero tilt.

IX. SYSTEMATICS

To evaluate detector systematics associated with the stability of detector over time, the data is compared across several deployments of the same source, collected months apart. The data for a central thorium source run are compared and a maximum of 5% bin-to-bin variations in the N_{hit}^p and time-residuals are identified. Other sources show similar sized maximum variations. This effect could be due to varying detector properties, such as the PMT gain and dark rate. To account for this, all figures comparing histograms of data and simulation in Secs. X and XI include a 5% systematic uncertainty added to each bin.

Differences in the reconstruction results over time are checked similarly. Variations in detector performance lead to a maximum difference in the SeedNDDestroy vertex reconstruction of the thorium source z position scan of 1.4 cm (mean position) and 0.35 cm (width) for the Cartesian coordinates. In presenting the reconstruction bias and width results with the thorium source, described in Sec. VI A, these maximum differences in performance are conservatively included as systematic uncertainties. The systematic is evaluated separately for each algorithm and is applied across all deployed source positions.

In addition to detector changes, there is an uncertainty in the knowledge of the true source position. This is particularly important to evaluate when comparing the reconstructed position distributions between data and simulation. As discussed in Sec. II the precision of the measurements of the various detector components, such as the height of the IV, leads to an uncertainty of 1.5 cm on z . Additionally, the measured tilt of the calibration deployment system, described in Sec. VIII G, is measured to have $\theta < 1^\circ$. Conservatively, the maximum ρ deviation is 0.9 cm at a tilt of 1° (negligible impact on z) and this systematic uncertainty is included in the vertex bias results for every deployed position (it has a negligible impact on the resolution). This systematic on the reconstructed position is summed in quadrature with the uncertainty assigned from the detector stability analysis and treated separately for each Cartesian coordinate and for each reconstruction algorithm.

In the direction reconstruction results presented in Sec. VI B, evaluation of the reconstruction is performed for when the source pointed directly downward and when $\theta = 90^\circ$. In the latter case, only a rough estimate of the ϕ direction was possible at the time of deploying the source,

accurate to approximately $\pm 5^\circ$. To account for this, the simulations are run scanning the value of ϕ in that range. The ϕ that produces the best data to simulation agreement is selected and a systematic uncertainty is generated based on the simulation with maximum difference in $\alpha_{1\sigma}$ from that value. In all cases, an additional systematic is estimated by adjusting the fit function to a single exponential, removing the Gaussian convolution. The difference between the measured $\alpha_{1\sigma}$ for the two models yields a systematic on the selected model.

In the following sections, systematic uncertainties are included on the simulation model when comparing data to simulation. Calculations of χ^2/NDF between the data and simulation include the systematic uncertainties, added in quadrature with the statistical uncertainty.

X. MODEL VALIDATION

The Eos calibration methods utilized the laserball source, with positions distributed across the central axis, to generate PMT cable delay, pulse-shape and charge calibration constants for the simulation. Afterward, a central thorium run is used to tune PMT efficiency corrections that are used to improve the N_{hit}^p data and simulation agreement. With these calibrated parameters, we can perform checks that the RAT-PAC-2 model is predictive by testing it using the off-center thorium source data as well as data with the other calibration sources. This section focuses on validating the time-residual and N_{hit}^p agreement between data and simulation, as these are the primary inputs into the reconstruction frameworks. While not as integral for the Eos reconstruction algorithms, the data and simulation agreement for various dichroicon data is also presented. The modeling of the dichroicons is of particular interest for future large-scale hybrid detector that may utilize these instruments for Cherenkov and scintillation separation.

A. Time residuals

The time residuals, described in Sec. VI A, are a critical input into the reconstruction frameworks. Figure 14 shows a histogram of t_{res} for the centrally deployed thorium source (left) and a centrally deployed, downward-pointing ^{90}Sr directional source (right) for both data and simulation. In this figure, the time-residuals are calculated using the reconstructed vertex from SeedNDDestroy (using the output from the other fit methods produces very similar results). The width of the residuals in the thorium and directional source data provide useful benchmarks. From a Gaussian fit to the peak, between -0.5 and 0.5 ns, widths (σ) of 278 ± 2 ps and 276 ± 2 ps are found in the thorium data and simulation, respectively (fit uncertainties only). Similarly, values of 247 ± 2 ps and 254 ± 4 ps are found in the directional source data and simulation. As the time-residuals around the prompt peak are the primary input for several of the likelihood-based reconstruction frameworks,




FIG. 14. Left: the time-residuals for the central thorium source data and MC, using SeedNDestroy to reconstruct the event vertex. Right: the same for the centrally deployed, downward-pointing ^{90}Sr directional source.

the good agreement between data and simulation indicated by these studies is an important check of the modeling.

B. N_{hit}

As described in Sec. VIII F, a central thorium run was used to tune an efficiency correction factor that improves the data and simulation agreement in $N_{\text{hit}}^{\text{p}}$. To validate the data and simulation agreement as a function of event position, we used the data from the thorium scan across the detector, in which data were collected in steps of 10 cm. Figure 15 shows the mean $N_{\text{hit}}^{\text{p}}$ value in the data and simulation, demonstrating agreement across a large range of event positions. This further verifies that the data and simulation agreement extends across a wide range of z -positions and demonstrates the validity of the optical modeling and efficiency tuning.

As further verification of the calibration, the data and MC agreement for $N_{\text{hit}}^{\text{p}}$ is checked for several other sources. This is shown for the selected Michel electrons (top right in Fig. 15), the centrally deployed 30 mm ^{90}Sr directional source, pointed both downward and sideways (central row of Fig. 15), and for the tagged prompt and delayed events from a centrally deployed AmBe source (bottom row in 15). For the Michel electron comparison, the $\chi^2/\text{NDF} = 11.6/25$. For the downward and sideways pointing directional sources, values of $\chi^2/\text{NDF} = 16.8/17$ and $\chi^2/\text{NDF} = 25.5/17$ are found respectively. For the AmBe prompt and delayed events, values of $\chi^2/\text{NDF} = 81.8/17$ and $\chi^2/\text{NDF} = 20.7/11$ are found, respectively. In general, the data and simulation agreement is very good across all sources, with the poorest agreement identified in the AmBe prompt events, due to a slight difference in the shape of the $N_{\text{hit}}^{\text{p}}$ distribution.

C. Dichroicon response

In situ measurements of the dichroicons were performed using the laserball source and compared to simulation. By

varying the laser wavelength and deployment height, the dataset effectively scans over the dichroicons' response both as a function of the photon wavelength and incidence angle. The collected data are compared with laserball simulations performed using RAT-PAC-2, as described in Sec. III. Spectrophotometry measurements from Ref. [45] are used to model the dichroic filters in the simulation. The Monte Carlo dataset is generated with the laserball positioned at all heights where data have been taken in the detector, and with identical run conditions as data collection. Rather than simulating monoenergetic photons at the specified laser wavelengths, a full wavelength range between 375 nm to 550 nm is simulated (in steps of 1 nm) to account for the finite wavelength spreads of the lasers. To account for laser intensity variations, occupancies of dichroicon PMTs are normalized by the average occupancy of the bottom nondichroic R14688-100 PMTs in the same run. This is referred to as the “normalized $N_{\text{hit}}^{\text{p}}$ ”.

Figure 16 shows the normalized $N_{\text{hit}}^{\text{p}}$ of the dichroicon PMTs in both data and MC for laserball deployments in the upper half of the detector. Below the cut-on wavelength of 450 nm, the absorbing long-pass filter that covers the PMT absorbs the photons, resulting in effectively zero hits. This is modeled well in the simulation. Around the cut-on wavelength, the optics of the dichroic filters are complicated, with transmission and reflection that depend strongly on the precise wavelength and incidence angle. The laser at 442 nm probes this transition region and shows good agreement between the data and simulation. This laser has the broadest wavelength distribution, which is indicated by the horizontal uncertainty. Above the cut-on wavelength, the absorbing filter becomes transparent, and the dichroic filters become reflective, allowing the dichroicon to act as a concentrator, increasing the effective coverage area of the PMTs. This results in these PMTs seeing more hits than the PMTs without dichroicons, as indicated by normalized occupancies above unity. At 515 nm the simulation slightly under-predicts the data at the center of the detector, and the




FIG. 15. Top left: the mean N_{hit}^p as a function of the source z position for the thorium data and simulation. The uncertainties are taken as the standard error on the mean. Top right: the N_{hit}^p for tagged Michel electrons compared for data and simulation. Center: the N_{hit}^p distributions for a central 30 mm ^{90}Sr directional run pointed (left) at 90 degrees (toward the barrel PMTs) and (right) directly downward. Bottom: the prompt (left) and delayed (right) event N_{hit}^p distributions from a central AmBe source run.

agreement between data and simulation improves at higher z . A $\chi^2/\text{NDF} = 9.8/7$ is calculated by summing over the seven deployment positions with the 515 nm laser. The agreement between the data and simulation gets worse at source positions below the center and this difference is understood to be caused by the dependence of the filters on incidence angle. There is ongoing work to better

understand and further calibrate the angular response of the dichroicon. It should be noted that in a large-scale detector, the distance between the events of interest and the PMTs is typically several meters, where the agreement between the data and MC for the dichroicons is best.

Additionally, the Michel electrons can be used to understand the dichroicon response to Cherenkov light, which




FIG. 16. Left: data and simulation comparison for the occupancy of dichroicon PMTs at different wavelengths and deployment heights. The occupancy is normalized by the average occupancy of 8 inch PMTs without dichroicons at the bottom PMT array for each run. The horizontal error bars are set based on the width (σ) of the laser emission distribution (which is described well as a Gaussian with a mean at the indicated wavelength). Right: the data and simulation comparison of the mean dichroicons N_{hit}^p as a function of the reconstructed z direction of the tagged Michel electrons.

spans the full wavelength regime probed by the lasers. The total occupancy of the PMTs with dichroicons in front of them is plotted as a function of the Mimir reconstructed z -direction ($\hat{d}_{recon} \cdot \hat{z}$) of the Michel electron in Fig. 16 (right). A value of $\chi^2/\text{NDF} = 11.4/24$ is calculated when comparing the data and simulation. As the Michel electrons point more directly downward, the average N_{hit}^p of the dichroicons increases.

XI. RECONSTRUCTION RESULTS

A major goal of the Eos water fill reconstruction results is to validate the methodology prior to collecting data with target materials that produce scintillation light. Performance metrics (e.g., bias and width distributions) for the reconstruction methods are shown to compare the performance of different algorithms and estimate the level of the data and simulation agreement. The capacity of the reconstruction is highly specific to the Eos detector, which is both small and asymmetric. As such, these results provide validation for future phases of Eos and for future simulation studies of large-scale hybrid detectors.

The reconstruction objectives for Eos are based on what has been achieved previously in large water Cherenkov detectors, such as SNO, SNO+, and Super-Kamiokande. These experiments find vertex bias and resolution systematics between 0.5 to 10 cm [46–48]. The systematics on the vertex bias and resolution tended to be subdominant in the total uncertainty for solar neutrino measurements presented in these references. Using the thorium source, Eos is evaluating vertex reconstruction at energies lower than any of these previous detectors. For this manuscript the goal for the agreement between data and MC in both the vertex bias and resolutions is to be within 3 cm, as characterized using the thorium source. This objective is

based on the experience of these previous water Cherenkov detectors, accounting for Eos’s unique, asymmetric detector configuration and low energy detection capabilities. Additionally, in principle any data and simulation differences could be either corrected for or taken as a systematic for any future Eos analysis that goes beyond evaluating detector performance.

The goals for the direction reconstruction, characterized using the directional sources, are determined similarly. For this manuscript, the goal is to provide data and simulation agreement in $\alpha_{1\sigma}$ to within 0.1 across various source z positions. This is similar to the size of the systematic on the angular resolution in Ref. [47], which resulted in a roughly 2% systematic on the ^8B flux measurement in a water Cherenkov detector.

A. Vertex

The event vertex is reconstructed using the methods described in Secs. VIA and VIC: the quad fitter, SeedNDestroy, Mimir and HitMan. Figure 17 shows the vertex reconstruction results in x , y , z , and ρ^2 for central thorium source data. In general, the quad fitter has the largest bias and resolution, as expected due to the simplicity of the algorithm. The other three fitters perform similarly, with HitMan slightly outperforming the likelihood-based fitters when reconstructing the z position. The reconstruction of the event z position is the most challenging as it is significantly impacted by the Eos top-bottom detector asymmetries.

For each of the algorithms, histograms of the reconstructed Cartesian coordinates are fit with a Gaussian around the peak, using a range of ± 20 cm. The bias (difference from the true source position) is taken from the mean of the fit and the width (σ) of the fit provides a




FIG. 17. Reconstructed x , y , z , and ρ^2/ρ_{IV}^2 distributions from the four vertex fitters applied to central thorium run data in Eos. The uncertainties are statistical only.

metric that estimates the vertex resolution. As is clear in Fig. 17, the distributions are not perfectly Gaussian, and the fit does not capture some of the features in the tails. This method is intended to demonstrate the performance for the bulk of the events without focusing on outliers. It should be emphasized that the thorium calibration source produces γ -rays, and thus, the reconstruction resolution convolves effects from the Compton scattering. As such, the goal here is to show the relative comparison between the algorithms, as well as the overall agreement between data and simulation.

Figure 18 shows the bias and resolution results for three fitters as a function of the source z -position (the quad fitter has been excluded, as it is used primarily to seed the other algorithms). The x and y bias from the true position for each of the fitters is around 1–2 cm, and does not depend strongly on the source z -position. The z bias ranges from 3–10 cm, is strongest in HitMan, and increases as the source moves higher in the detector. This is expected, as the PMT coverage toward the bottom of the detector is higher than at the top, and the simulation accurately predicts this effect. The systematic uncertainties on the bias for each coordinate are on the order of 2 cm, as described in Sec. IX. The χ^2/NDF comparisons between the data and simulation are

provided for each reconstruction algorithm for both the bias and resolution in Table III. In several cases, a very small value of the χ^2/NDF is indicative that the systematic uncertainties have been overestimated. This is expected in some cases as the systematics were estimated based on conservative assumptions.

The x and y resolutions (σ) from the Gaussian fits are around 12–13 cm for all algorithms and the data and simulation agree within 1 cm. The z resolution has an understood dependence on the source z -position, and extends from roughly 12 cm near the bottom to 16 cm near the top of the detector. The data and simulation agreement is most clearly delineated in Fig. 19, which shows, for each algorithm, the difference between the data and simulation for the bias and resolution. In general, deviations between the data and simulation for both the bias and resolution are on the order of 1–2 cm for each algorithm. Interestingly, the simulation tends to slightly overpredict the resolution when compared to the data.

The AmBe source provides an additional method for testing the vertex reconstruction, with a different source geometry (and thus different source shadowing effects) and different energy γ -rays (both the prompt 4.4 MeV and delayed 2.2 MeV). In Fig. 20 the SeedNDestroy reconstructed




FIG. 18. The bias (left) and resolution (right) for the x (top), y (middle), and z (bottom) positions, for SeedNDDestroy, Mimir, and HitMan, as a function of the thorium source z -position. The values are taken from Gaussian fits performed over the histogram of the relevant reconstructed quantity. For Mimir and HitMan the source z position is shifted slightly from the true value on the x axis to make the figure more legible.

x , y , and z position distributions are compared between data and MC for the tagged prompt and delayed events from a centrally deployed AmBe source. The vertex resolution for the prompt/delayed events from the 4.4/2.2 MeV γ -rays is around 13/19 cm in the data. The data and simulation for these reconstructed distributions are compared and the

χ^2/NDF are 64.1/27 (prompt x), 66.4/27 (prompt y), 85.7/27 (prompt z), 47.8/27 (delayed x), 31.4/27 (delayed y), and 74.6/27 (delayed z). Similarly to the $N_{\text{hit}}^{\text{p}}$ distributions, there remain small mismodelings of the AmBe source in the detector that lead to relatively large χ^2/NDF values.

TABLE III. The χ^2/NDF values for the data and simulation comparisons presented in Fig. 18.

	SeedNDestroy	Mimir	HitMan
x bias	10.4/11	10.7/11	10.6/11
y bias	4.8/11	2.0/11	1.8/11
z bias	11.5/11	10.6/11	8.8/11
x resolution	2.7/11	1.8/11	17.5/11
y resolution	1.5/11	5.3/11	17.0/11
z resolution	16.1/11	14.3/11	10.0/11

The Δr distribution between the prompt and delayed events for the AmBe is shown in Fig. 21. Included is the background distribution from a dedicated run without the source in the detector, scaled to have the same livetime as the AmBe run. For this figure, the fiducial volume is relaxed to show the full background distribution. The Δr distribution exemplifies that, despite the surface deployment and water target, Eos can measure low-energy coincidences from neutrons with little background. The good data and simulation agreement indicates that while we observe small differences between the data and simulation in the individual Cartesian coordinates, overall the Δr distribution is modeled quite well ($\chi^2/\text{NDF} = 21.5/26$).

In Eos, the directional source data provide a test of the vertex reconstruction using electrons. Due to source shadowing, the β energy spectrum, and other effects, we must still rely on our simulation model to understand the resolutions. The figures for the directional source results are detailed in Appendix A. In these figures, we combine across several runs in order to generate enough statistics for the reconstruction tests. These include data and simulation comparisons for both the ^{90}Sr and ^{106}Ru source, the latter of which includes two source sizes (20 mm and 30 mm). For the ^{90}Sr source, the bias and resolution for both downward pointing and sideways pointing data are displayed. Overall, the x and y resolutions are less than 10 cm for the ^{106}Ru data

across all source z -positions. As in all cases, the resolution in z is more complicated and has a stronger dependence on the deployed source z -position.

As a useful metric, we provide the resolutions for simulated central, 2.0 MeV electrons in Eos for each of the algorithms in Table IV. For all algorithms other than the quad fitter, the expected resolution is less than 10 cm for central 2.0 MeV electrons.

The agreement achieved in the bias and resolutions across varying source type, particle type, source size, source position, and source direction (with energies spanning roughly 1 to 4.5 MeV) is typically around 1 cm and is less than 3 cm for all algorithms. These efforts provide additional verification of the Eos modeling, reconstruction software, and the model of the sources themselves. This is critical for future phases of Eos using various liquid scintillators as the target material. The evaluation of the difference between data and simulation in Eos allows us to understand the limitations of our predictions for these future phases and for future detectors, such as Theia.

B. Direction

The event direction is reconstructed using the methods described in Sec. VI B. Figure 22 compares the direction reconstruction results of α for the data and simulation, applying the four different direction reconstruction algorithms to the central 30 mm ^{106}Ru directional source pointed downward. Due to the self-triggering mechanism of the directional source, negligible background is identified, which would manifest as a flat component in $\hat{a}_{\text{source}} \cdot \hat{a}_{\text{recon}}$. All four algorithms perform similarly and show comparable levels of agreement between data and simulation. The calculated χ^2/NDF values for these data and simulation comparisons are 36.7/28 (average direction), 33.9/30 (Mimir), 34.3/29 (AERO), and 38.6/30 (HitMan). Based on these findings, the following directional




FIG. 19. The data minus the simulation for the x , y , and z bias and resolutions from the three vertex fitters, for the thorium source. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 20. The reconstructed position distributions, using SeedNDestroy, compared between data and simulation for a central AmBe source run. The left column shows the prompt event and the right column shows the delayed event.

source results focus on the agreement for a single reconstruction framework (AERO).

These histograms are fit with an exponential convolved with a Gaussian, described in Sec. VIB, to extract $\alpha_{1\sigma}$. In this manuscript, we present selected results for the directional sources: the downward pointing 30 mm ^{90}Sr and ^{106}Ru sources, the downward pointing 20 mm ^{106}Ru source, and the 90° sideways pointing 20 mm ^{106}Ru source. In general, consistent performance is identified with the 20 mm and 30 mm sources, so not all of the

collected data is presented here. Additionally, the source z position is limited to $[-30, +30]$ cm in these studies. Data collected closer to the bottom of the detector shows poorer performance that is not yet well modeled in the simulation, similar to the discussion about the dichroicon performance with the laserball data in Sec. XC. There is active, ongoing work to better understand the detector response for sources near the bottom of the detector. For the sideways pointing runs, an additional systematic is included (described in Sec. IX) based on the fact that




FIG. 21. The reconstructed Δr distribution using SeedNDestroy, compared between data and simulation for a central AmBe source run. The background distribution is taken from a dedicated run without the source in the detector and is scaled to the same livetime as the AmBe data.

no precise measurement of the ϕ direction of the source was possible at the time.

The results for $\alpha_{1\sigma}$ for these deployments as a function of the source z position are shown in Fig. 23. In general, the

TABLE IV. The x , y , and z resolution from a Gaussian fit to the associated histogram for simulated central, isotropically pointed, 2.0 MeV electrons. The uncertainties are from the fit and are statistical only.

Algorithm	x (cm)	y (cm)	z (cm)
Quad	10.9 ± 0.2	10.8 ± 0.2	10.7 ± 0.2
SeedNDestroy	8.9 ± 0.1	8.8 ± 0.1	9.7 ± 0.1
Mimir	8.8 ± 0.1	8.8 ± 0.1	9.4 ± 0.1
HitMan	7.8 ± 0.1	7.9 ± 0.1	8.3 ± 0.1

data and simulation agreement is good and the χ^2/NDF values for the 30 mm ^{90}Sr source (downward), 30 mm ^{106}Ru source (downward), 20 mm ^{106}Ru source (downward), and 30 mm ^{106}Ru source (sideways) are 1.8/7, 1.9/7, 2.1/7, and 1.7/7, respectively. Across these results we find that the agreement between the data and simulation is well within the goal of 0.1 on $\alpha_{1\sigma}$, regardless of the source type and direction.

The direction reconstruction can also be tested using Michel electrons, which provide a higher energy, distributed source in which to study the algorithms. Figure 24 shows $\hat{d}_{\text{recon}} \cdot \hat{d}_{\text{PMT}}$, where $\hat{d}_{\text{PMT}}$ is the vector pointing from the hit PMT to the reconstructed event position for the




FIG. 22. The α reconstruction results for data and simulation, from a central deployment of the 30 mm ^{106}Ru directional source, pointed downward. All four direction reconstruction algorithms are shown.




FIG. 23. The $\alpha_{1\sigma}$ reconstruction results using AERO for data and simulation for the downward pointing 30 mm ^{90}Sr directional source (top left), the downward pointing 30 mm ^{106}Ru directional source (top right), the downward pointing 20 mm ^{106}Ru directional source (bottom left), and the $\theta = 90^\circ$ sideways pointing 30 mm ^{90}Sr directional source (bottom right).

tagged Michel electron. This distribution is shown using both all PMTs (left) and only the PMTs behind the dichroicons (right). These results use Mimir for the reconstructed direction. The peak in both distributions around $\hat{d}_{\text{recon}} \cdot \hat{d}_{\text{PMT}} \sim 0.75$ is from the Cherenkov angle at $\sim 41^\circ$.

The slight offset in the peak from the Cherenkov angle, in both the data and simulation, in the dichroicon-only result is caused by geometric effects. Values of $\chi^2/\text{NDF} = 13.6/24$ and $\chi^2/\text{NDF} = 53.4/24$ are calculated for the $\hat{d}_{\text{recon}} \cdot \hat{d}_{\text{PMT}}$ results for the all PMTs and dichroicon PMTs




FIG. 24. The $\hat{d}_{\text{recon}} \cdot \hat{d}_{\text{PMT}}$ compared between data and simulation using all PMTs (left) and using only the PMTs behind the dichroicons (right). The dashed gray line denotes the Cherenkov emission angle.

respectively. The agreement between data and simulation is further verification of the detector modeling and direction reconstruction, which work well even in the 10's of MeV range for tagged Michel electrons.

XII. DISCUSSION

Eos plays a key role as the first large-scale demonstrator of many of the key technologies that enable future hybrid detectors. This manuscript marks the first use of Eos data and demonstrates the development and calibration of simulation models that can be used to make predictions and generate physics sensitivity estimates for Theia. These calibrated models will also be invaluable in the subsequent phases of Eos, which include the deployment of both water-based liquid scintillator (WbLS) [11] and pure liquid scintillator (such as LAB + PPO) [49], where the optical properties of the medium in the IV is less well understood.

Furthermore, the AmBe results show that Eos is able to use the deployed advanced instrumentation and sophisticated triggering to perform clear neutron detection, despite the near-surface deployment, the small detector size, and the water target, which produces only Cherenkov light. Neutron detection is a key component for the detection of reactor antineutrinos, and thus these results aid in understanding the potential to use this technology for nuclear reactor monitoring and other important physics topics that require the detection of antineutrinos or neutrons.

The Eos Collaboration is continuing to develop detector modeling and data analysis techniques. This includes actively developing multi-PE counting and timing estimators, utilizing charge in the reconstruction algorithms, improving the modeling of the AmBe source, analyzing additional calibration source data from sources not used for this paper (e.g., an external γ source), high energy (> 10 MeV) reconstruction methods, angular response calibrations of the dichroicons, and improvements to the incorporation of scintillation light into the reconstruction algorithms. This work is paired with ongoing efforts to understand physics sensitivities for a potential future deployment of Eos at the Spallation Neutron Source at Oak Ridge National Lab.

XIII. CONCLUSION

In this manuscript we describe the data collected with the Eos detector while it was filled with water, using a variety of deployed radioactive and optical sources. Laserball data are used to calibrate the PMT timing and charge, tune the model of the dichroicons, and investigate the trigger efficiency. Central thorium data are used to calibrate the PMT efficiencies. The results from these calibrations are input into detailed simulation models, built with the RAT-PAC-2 software. Using this framework, a variety of checks are performed against the data using the laserball, thorium, AmBe, and several directional sources. These analyses indicate good data and simulation agreement for the PMT

time-residuals and N_{hit}^p , across a wide range of source positions. Both likelihood- and ML-based reconstruction methods that utilize the PMT time-residuals, are used to reconstruct the event vertex and direction. The outputs of these reconstruction methods are compared between data and simulation for the thorium, directional, and AmBe sources. Agreement between data and simulation to within the stated goals for Eos is achieved for the reconstructed parameters. This is a powerful statement as these tools are being used to extrapolate to much larger detectors and to evaluate the physics sensitivities of future detectors. Additionally, the detailed calibrations, detector models, and calibration source models developed for the water phase are critical to the success of future phases of Eos in which less well-understood target materials will be evaluated. The results from this manuscript will aid in the future demonstration of Cherenkov and scintillation separation in a tonne-scale hybrid detector and are a critical step toward evaluating the performance of many key elements of hybrid detector technology.

ACKNOWLEDGMENTS

Work conducted at Lawrence Berkeley National Laboratory was performed under the auspices of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231. The work conducted at Lawrence Livermore National Laboratory was supported by the U.S. Department of Energy under Contract No. DE-AC52-07NA27344, release number LLNL-JRNL-2020089. The work conducted at Brookhaven National Laboratory was supported by the U.S. Department of Energy under Contract No. DE-AC02-98CH10886. The work conducted at Oak Ridge National Laboratory was supported by the U.S. Department of Energy under Contract No. DE-AC05-00OR22725. The work conducted at Oxford was supported by the Royal Society. The project was funded by the U.S. Department of Energy, National Nuclear Security Administration, Office of Defense Nuclear Nonproliferation Research and Development (DNN R&D). This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of High Energy Physics, under Award No. DE-SC0018974. This work is supported as part of the Growth in Research Opportunities With Traineeship in High energy physics at Minority Serving Institutions (GROWTH-MSI) program funded by the DOE Office of Sciences High Energy and Nuclear Physics, Grant No. DE-SC0023725. This research used the Savio computational cluster resource provided by the Berkeley Research Computing program at the University of California, Berkeley (supported by the UC Berkeley Chancellor, Vice Chancellor for Research, and Chief Information Officer).

DATA AVAILABILITY

The Eos raw data was collected as hdf5 files and access can be provided upon reasonable request. The Eos

processed data and simulated data are stored in ROOT format, and access can be provided upon reasonable request. All data and simulation is stored on the Savio computing cluster at the University of California, Berkeley. All data processing and simulation generation was performed using the open source RAT-PAC-2 framework.

APPENDIX A: DIRECTIONAL SOURCE VERTEX RECONSTRUCTION

In this section we show the vertex reconstruction results for SeedNDEstroy, Mimir, and HitMan for various directional source data. The reconstructed x , y , and z bias and resolutions are shown in Figs. 25–28 for the downward




FIG. 25. The bias (left) and resolution (right) for the x (top), y (middle), and z (bottom) positions, for SeedNDEstroy, Mimir, and HitMan, for the downward pointing 30 mm ^{90}Sr directional source. The values are taken from Gaussian fits performed over the histogram of the relevant reconstructed quantity. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 26. The bias (left) and resolution (right) for the x (top), y (middle), and z (bottom) positions, for SeedNDestroy, Mimir, and HitMan, for the downward pointing 30 mm ^{106}Ru directional source. The values are taken from Gaussian fits performed over the histogram of the relevant reconstructed quantity. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 27. The bias (left) and resolution (right) for the x (top), y (middle), and z (bottom) positions, for SeedNDestroy, Mimir, and HitMan, for the downward pointing 20 mm ^{106}Ru directional source. The values are taken from Gaussian fits performed over the histogram of the relevant reconstructed quantity. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 28. The bias (left) and resolution (right) for the x (top), y (middle), and z (bottom) positions, for SeedNDestroy, Mimir, and HitMan, for the 90° sideways pointing 30 mm ^{90}Sr directional source. The values are taken from Gaussian fits performed over the histogram of the relevant reconstructed quantity. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 29. The data minus the simulation for the x , y , and z bias and resolutions from the three vertex fitters, for the downward pointing 30 mm ^{90}Sr directional source. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 30. The data minus the simulation for the x , y , and z bias and resolutions from the three vertex fitters, for the downward pointing 30 mm ^{106}Ru directional source. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 31. The data minus the simulation for the x , y , and z bias and resolutions from the three vertex fitters, for the downward pointing 20 mm ^{106}Ru directional source. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.




FIG. 32. The data minus the simulation for the x , y , and z bias and resolutions from the three vertex fitters, for the 90° sideways pointing 30 mm ^{90}Sr directional source. For Mimir and HitMan the source z positions are shifted slightly from the true value on the x axis to make the figure more legible.

pointing 30 mm ^{90}Sr , the downward pointing 30 mm ^{106}Ru , the downward pointing 20 mm ^{106}Ru , and the sideways pointing 30 mm ^{90}Sr respectively. The differences between the data and simulations are shown in Figs. 29–32 respectively. For the sideways pointing data, the azimuth of the source is fixed to 175 degrees in the (x, y) plane and the polar angle (θ) is set to 90 degrees.

As with the thorium data described in Sec. XI A, the data and simulation tend to agree to within 1–2 cm for both the bias and resolution. Overall, the x and y resolutions (σ) are less than 10 cm. The z resolution is around 8–12 cm, and depends more strongly on the source z -position. The resolutions are slightly broader for the sideways pointing source (~ 8 –12 cm) than for the downward pointing source (~ 8 –10 cm) due to the better coverage at the bottom of the detector than for the barrel. Additionally, the resolution is slightly broader for the ^{90}Sr data than the ^{106}Ru data, due to the slightly lower end point energy.

APPENDIX B: MUON LIFETIME

This section calculates the expected muon lifetime in water. Negative muons in water can undergo nuclear capture on oxygen, reducing their effective lifetime relative

to the vacuum value. The corresponding decay rates of positive and negative muons are

$$\frac{1}{\tau_{\mu^+}} = R_{\text{vac}}, \quad \frac{1}{\tau_{\mu^-}} = R_{\text{vac}} + R_{\text{oxy}}, \quad (\text{B1})$$

where R_{vac} is the vacuum muon decay rate and R_{oxy} is the capture rate on oxygen. With a μ^+/μ^- ratio, r , the total probability that a muon at rest in water decays is

$$\begin{aligned} P(\text{decay}) &= P(\text{decay}|\mu^+)P(\mu^+) + P(\text{decay}|\mu^-)P(\mu^-) \\ &= \frac{r}{1+r} + \frac{R_{\text{vac}}}{R_{\text{vac}} + R_{\text{oxy}}} \frac{1}{1+r}. \end{aligned} \quad (\text{B2})$$

The effective muon lifetime measured by tagging decays from both μ^+ and μ^- is the posterior-weighted average of the μ^+ and μ^- lifetimes:

$$\tau_{\text{tot}} = \tau_{\mu^+}P(\mu^+|\text{decay}) + \tau_{\mu^-}P(\mu^-|\text{decay}). \quad (\text{B3})$$

Using the experimental results in [50,51], the expected muon lifetime in water is $2.04 \pm 0.01 \mu\text{s}$.

-
- [1] Y. Fukuda *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. Lett.* **81**, 1562 (1998).
 - [2] Q. R. Ahmad *et al.* (SNO Collaboration), *Phys. Rev. Lett.* **89**, 011301 (2002).
 - [3] R. Bonventre and G. D. Orebi Gann, *Eur. Phys. J. C* **78**, 435 (2018).

- [4] S. Andringa *et al.* (SNO+ Collaboration), *Adv. High Energy Phys.* **2016**, 6194250 (2016).
- [5] M. Agostini *et al.* (BOREXINO Collaboration), *Phys. Rev. Lett.* **128**, 091803 (2022).
- [6] A. Allega *et al.* (SNO+ Collaboration), *Phys. Rev. D* **109**, 072002 (2024).

- [7] M. Askins *et al.* (Theia Collaboration), *Eur. Phys. J. C* **80**, 416 (2020).
- [8] J. Sawatzki, M. Wurm, and D. Kresse, *Phys. Rev. D* **103**, 023021 (2021).
- [9] M. Askins *et al.* (WATCHMAN Collaboration), *arXiv:1502.01132*.
- [10] A. Bat, E. Tiras, V. Fischer, and M. Kamislioglu, *Eur. Phys. J. C* **82**, 734 (2022).
- [11] M. Yeh, S. Hans, W. Beriguete, R. Rosero, L. Hu, R. L. Hahn, M. V. Diwan, D. E. Jaffe, S. H. Kettell, and L. Littenberg, *Nucl. Instrum. Methods Phys. Res., Sect. A* **660**, 51 (2011).
- [12] S. D. Biller, E. J. Leming, and J. L. Paton, *Nucl. Instrum. Methods Phys. Res., Sect. A* **972**, 164106 (2020).
- [13] Z. Guo, M. Yeh, R. Zhang, D.-W. Cao, M. Qi, Z. Wang, and S. Chen, *Astropart. Phys.* **109**, 33 (2019).
- [14] H. T. J. Steiger *et al.*, *J. Instrum.* **19**, P09015 (2024).
- [15] A. V. Lyashenko *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **958**, 162834 (2020).
- [16] T. Kaptanoglu, E. J. Callaghan, M. Yeh, and G. D. O. Gann, *Eur. Phys. J. C* **82**, 169 (2022).
- [17] T. Kaptanoglu, M. Luo, and J. Klein, *J. Instrum.* **14**, T05001 (2018).
- [18] T. Kaptanoglu, M. Luo, B. Land, A. Bacon, and J. Klein, *Phys. Rev. D* **101**, 072002 (2020).
- [19] A. Elagin, H. Frisch, B. Naranjo, J. Ouellet, L. Winslow, and T. Wongjirad, *Nucl. Instrum. Methods Phys. Res., Sect. A* **849**, 102 (2017).
- [20] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), *Phys. Rev. Lett.* **135**, 171804 (2025).
- [21] T. Anderson *et al.*, *J. Instrum.* **18**, P02009 (2022).
- [22] Hamamatsu, Photomultiplier tube R14688/R14688-100, https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/R14688-R14688-100_TPMH1385E.pdf (2022).
- [23] T. Kaptanoglu, A. Rincon, M. Duce, S. Kaplan, J. Koplowitz, S. Lynch, H. J. Ryoo, and G. Orebi Gann, *J. Instrum.* **19**, P02032 (2023).
- [24] J. Brack *et al.*, *Nucl. Instrum. Methods Phys. Res., Sect. A* **712**, 162 (2013).
- [25] Hamamatsu, Large area photocathode photomultiplier tubes, https://www.hamamatsu.com/content/dam/hamamatsu-photonics/sites/documents/99_SALES_LIBRARY/etd/LARGE_AREA_PMT_TPMH1376E.pdf (2022).
- [26] T. Anderson *et al.* (Eos Collaboration), *arXiv:2607.20285*.
- [27] CAEN, CAEN V1730, <https://www.caen.it/products/v1730/> (2025).
- [28] S. Seibert *et al.*, RatPac-two user's guide, <https://ratpac.readthedocs.io/en/latest/> (2025).
- [29] S. Agostinelli *et al.* (GEANT4 Collaboration), *Nucl. Instrum. Methods Phys. Res., Sect. A* **506**, 250 (2003).
- [30] GEANT4 Collaboration, Guide for physics lists, <https://geant4-userdoc.web.cern.ch/UsersGuides/PhysicsListGuide/BackupVersions/V11.1b/html/index.html> (2023).
- [31] W. Warburton and W. Hennig, *IEEE Trans. Nucl. Sci.* **64**, 2938 (2017).
- [32] W. Jiang, G. Huang, Z. Liu, W. Luo, L. Wen, and J. Luo, *Eur. Phys. J. C* **85**, 69 (2025).
- [33] M. Akashi-Ronquest *et al.*, *Astropart. Phys.* **65**, 40 (2015).
- [34] J. Tang, T. Xiao, X. Tang, and Y. Huang, *J. Instrum.* **20**, P03019 (2024).
- [35] D. Xu *et al.*, *J. Instrum.* **17**, P06040 (2021).
- [36] M. Daimon and A. Masumura, *Appl. Opt.* **46**, 3811 (2007).
- [37] I. Coulter, Modelling and reconstruction of events in SNO+ related to future searches for lepton and baryon number violation, Ph.D. thesis, University of Oxford, Oxford, 2013.
- [38] S. G. Johnson, The NLOpt nonlinear-optimization package, <https://github.com/stevengj/nlopt> (2007).
- [39] T. H. Rowan, Functional stability analysis of numerical algorithms, Ph.D. thesis, Department of Computer Science, University of Texas at Austin, Austin, TX, 1990.
- [40] F. James and M. Roos, *Comput. Phys. Commun.* **10**, 343 (1975).
- [41] K. M. Gorski, E. Hivon, A. J. Banday, B. D. Wandelt, F. K. Hansen, M. Reinecke, and M. Bartelmann, *Astrophys. J.* **622**, 759 (2005).
- [42] P. Eller, A. T. Fienberg, J. Weldert, G. Wendel, S. Böser, and D. F. Cowen, *Nucl. Instrum. Methods Phys. Res., Sect. A* **1048**, 168011 (2023).
- [43] J. Hermans, V. Begy, and G. Louppe, in *International Conference on Machine Learning* (PMLR, Cambridge, MA, 2020), pp. 4239–4248.
- [44] D. Misra, *arXiv:1908.08681*.
- [45] A. Bacon, B. Harris, and J. Klein, *J. Instrum.* **20**, P09012 (2025).
- [46] B. Aharmim *et al.* (SNO Collaboration), *Phys. Rev. C* **81**, 055504 (2010).
- [47] A. Allega *et al.* (SNO+ Collaboration), *Phys. Rev. D* **110**, 122003 (2024).
- [48] K. Abe *et al.* (Super-Kamiokande Collaboration), *Phys. Rev. D* **94**, 052010 (2016).
- [49] M. R. Anderson *et al.* (SNO+ Collaboration), *J. Instrum.* **16**, P05009 (2020).
- [50] T. Suzuki, D. F. Measday, and J. P. Roalsvig, *Phys. Rev. C* **35**, 2212 (1987).
- [51] T. H. Burnett, L. J. LaMay, G. E. Masek, T. Maung, E. S. Miller, H. Ruderman, and W. Vernon, *Phys. Rev. Lett.* **30**, 937 (1973).