Axionlike particle search with a hybrid Cherenkov-scintillation detector

A. A. Aguilar-Arevalo,[8] S. Biedron,[12] J. Boissevain,[1] M. Borrego,[6] L. Bugel,[7,†] M. Chavez-Estrada,[8] J. M. Conrad,[7] R. L. Cooper,[6,10] J. R. Distel,[6] J. C. D'Olivo,[8] E. Dunton,[6] B. Dutta,[11] D. E. Fields,[9] M. Gold,[9] E. Guardincerri,[6] E. C. Huang,[6] N. Kamp,[7] D. Kim,[13] K. Knickerbocker,[6] W. C. Louis,[6] C. F. Macias-Acevedo,[8] R. Mahapatra,[11] J. Mezzetti,[5] J. Mirabal,[6] M. J. Mocko,[6] D. A. Newmark,[7,*] P. deNiverville,[6] C. O'Connor,[5] V. Pandey,[14] D. Poulson,[6] H. Ray,[5] E. Renner,[6] T. J. Schaub,[9] A. Schneider,[11] M. H. Shaevitz,[2] D. Smith,[4] W. Sondheim,[6] A. M. Szelc,[3] C. Taylor,[6] A. Thompson,[15] W. H. Thompson,[6] R. T. Thornton,[6] M. Tripathi,[5] R. Van Berg,[1] R. G. Van de Water,[6] and J. Zettlemoyer[6]

(CCM Collaboration)

[1]*Bartoszek Engineering, Aurora, Illinois 60506, USA*
[2]*Columbia University, New York, New York 10027, USA*
[3]*University of Edinburgh, Edinburgh, United Kingdom*
[4]*Embry–Riddle Aeronautical University, Prescott, Arizona 86301, USA*
[5]*University of Florida, Gainesville, Florida 32611, USA*
[6]*Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA*
[7]*Massachusetts Institute of Technology, Cambridge, Massachusetts 02139, USA*
[8]*Universidad Nacional Autónoma de México, CDMX 04510, México*
[9]*University of New Mexico, Albuquerque, New Mexico 87131, USA*
[10]*New Mexico State University, Las Cruces, New Mexico 88003, USA*
[11]*Texas A&M University, College Station, Texas 77843, USA*
[12]*Element Aero, San Leandro, California 94577, USA*
[13]*University of South Dakota, Vermillion, South Dakota 57069, USA*
[14]*Fermi National Accelerator Laboratory, Batavia, Illinois 60510, USA*
[15]*Northwestern University, Evanston, Illinois 60208, USA*

(Received 17 August 2026; accepted 8 October 2026; published 27 October 2026)

This analysis presents the first proof-of-concept search for axionlike particles using a hybrid Cherenkov-scintillation detector at a beam-dump facility. The work is based on the Coherent CAPTAIN-Mills (CCM) experiment, a 10-ton liquid argon light collection detector located at Los Alamos National Laboratory. The CCM200 detector is instrumented with 200 photomultiplier tubes, providing approximately 50% photocathode coverage. To enable optical discrimination, 80% of the photomultiplier tubes are coated with a wavelength-shifting material while the remaining 20% are left uncoated. This configuration, combined with nanosecond-scale timing resolution, provides sensitivity to prompt Cherenkov radiation while maintaining efficient detection of liquid argon scintillation light [Phys. Rev. Lett. **135**, 171804 (2025); [Phys. Rev. D **112**, 072010 (2025)]. This analysis constructs four observables that exploit the Cherenkov emission and event topology expected from axionlike-particle-induced electromagnetic interactions. These observables are combined into a likelihood-ratio classifier that provides powerful rejection of steady-state backgrounds. No statistically significant excess above the expected background prediction is observed for 10^{-3} MeV $< m_a <$ 10 MeV. Nevertheless, the improved background rejection enabled by hybrid Cherenkov-scintillation detection allows this analysis to surpass the sensitivity of the previous CCM120 search [Phys. Rev. D **107**, 095036 (2023)] despite less exposure and demonstrate the physics potential for hybrid Cherenkov-scintillation detectors.

DOI: 10.1103/ztsv-s7ts

*Contact author: dnewmark@mit.edu
†Deceased

Published by the American Physical Society under the terms of the Creative Commons Attribution 4.0 International license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI. Funded by SCOAP³.

I. INTRODUCTION

The axion arises from the spontaneous breaking of the global Peccei-Quinn $U(1)_{PQ}$ symmetry [1], proposed to explain the observed conservation of charge-parity (CP) symmetry in the strong interaction by dynamically solving the strong-CP problem [2,3]. The resulting QCD

axion [4–7] is accompanied by a broader class of axionlike particles (ALPs), which arise in many extensions of the Standard Model. In addition to solving the strong-*CP* problem, light axions and ALPs are also well-motivated dark matter and dark sector candidates [8].

Although a broad experimental program is dedicated to searching for very light axions and ALPs (ADMX [9], CAST [10], IAXO [11], HAYSTAC [12,13], and DMRadio [14] among others), this work targets higher-mass ALP candidates. Accelerator- and reactor-based experiments (NA62 [15], NA64 [16,17], SeaQuest [18], DarkQuest [19], NEON [20], and MINER [21] among others) offer a complementary avenue for exploring MeV to GeV mass ALP parameter space, which remains comparatively less constrained. ALPs in this mass range could still address the strong-*CP* problem and provide a window into a more extensive dark sector.

To search for these rare interactions, optical scintillation detectors are a well-established technology that offer large target masses, high light yield, and excellent energy resolution [22–25]. More recently, hybrid optical detectors capable of simultaneously measuring both prompt Cherenkov emission and delayed scintillation light have emerged as a promising approach for improving event reconstruction and background rejection [26–29]. Because Cherenkov radiation is emitted at a characteristic angle, reconstructing both signals provides additional information for particle identification and topology beyond that available from scintillation light alone.

The Coherent CAPTAIN-Mills (CCM) experiment at the Lujan Spallation Facility of the Los Alamos Neutron Science Center (LANSCE) employs a liquid argon optical detector designed to leverage scintillation and Cherenkov light for neutrino and beyond the Standard Model physics searches [30–33]. A previous search for ALPs was performed with the prototype CCM120 detector [34], using an exposure of 1.79×10^{21} protons on target (POT). This prototype detector was instrumented with 120 photomultiplier tubes (PMTs) and achieved a timing resolution of $\mathcal{O}(10\ \text{ns})$, which limited its ability to separate the prompt Cherenkov signal from the slower scintillation emission. Since this separation requires timing resolution approaching $\mathcal{O}(1\ \text{ns})$, the CCM120 analysis relied primarily on timing cuts within the 290 ns beam spill window to suppress steady-state backgrounds. Despite this limitation, the analysis demonstrated the potential of accelerator-based neutrino facilities to probe previously unexplored ALP parameter space.

Since the completion of the CCM120 program, both the detector and data processing techniques have undergone substantial upgrades. The current CCM200 detector features 200 PMTs, an external muon-tagging veto system, and improved data acquisition capabilities enabling more extensive calibration studies. A key feature of the CCM200 design is the combination of wavelength-shifting coated and uncoated PMTs to enable separation of scintillation and Cherenkov light. Specifically, 80% of the PMTs and the detector walls are evaporatively coated with the wavelength shifter tetraphenyl butadiene (TPB), while the remaining 20% of PMTs are uncoated. The TPB converts the 128 nm liquid argon scintillation light into visible wavelengths detectable by the PMTs, while the uncoated PMTs retain enhanced sensitivity to the prompt visible component of the broad-spectrum Cherenkov radiation.

In addition to these detector improvements, Ref. [35] presents new calibration and data processing techniques developed for the CCM200 program. These include pulse unfolding to reconstruct single photoelectron hits with 2 ns timing resolution and calibration studies using a ^{22}Na radioactive source producing $\mathcal{O}(1\ \text{MeV})$ gamma rays. The combination of wavelength-selective PMTs and improved timing resolution enabled event-by-event separation of scintillation and Cherenkov light from sub-MeV electrons [26]. This capability establishes CCM200 as a hybrid optical detector that combines the high light yield of scintillation emission with the directional and particle identification information provided by Cherenkov radiation.

The ability to identify Cherenkov light provides a new handle for distinguishing signal-like electromagnetic interactions from beam-related and environmental backgrounds that are difficult to separate using scintillation information alone. The distinct spatial and temporal signatures of Cherenkov and scintillation light enable the construction of new discriminating observables, which are incorporated in this analysis through the event selection. Building upon the CCM200 data processing and optical model calibration framework established in Refs. [26,35], this analysis improves background rejection and enhances sensitivity to ALP-induced signals compared to the CCM120 result. It provides a proof-of-principle demonstration of hybrid Cherenkov-scintillation detection in a beam-dump experiment and represents the first ALP search performed with a hybrid optical detector at an accelerator-based beam-dump facility. Further details of this analysis are presented in Ref. [36].

In this work, we revisit the ALP search using data collected during approximately eight weeks in 2022 with the CCM200 detector, corresponding to a total of 1.23×10^{21} POT. This analysis considers ALP masses $10^{-3}\ \text{MeV} < m_a < 10\ \text{MeV}$ and photon coupling strengths $10^{-6}\ \text{GeV}^{-1} < g_{a\gamma} < 10^{-3}\ \text{GeV}^{-1}$. During data collection, PMT waveforms were recorded in a 16 μs window around each beam spill, allowing for data-driven background characterization. The data collected before the beam spill are used to characterize the steady-state background, referred to as the "prebeam" sample. These events are dominated by low-energy activity from electronics noise, ambient radioactivity in the detector hall, and cosmogenic backgrounds.

In addition to the prebeam sample, this work uses events from the neutron-dominated time region as a data-driven control sample to validate the electromagnetic event selection. These neutron-induced events are generally higher in energy than those in the prebeam sample, providing a complementary test of the selection at higher energies. Following the proton beam interaction with the tungsten target, prompt neutrinos and near-speed-of-light beyond the Standard Model particles (such as ALPs) reach the detector with minimal delay, whereas the larger neutron flux arrives approximately 200 ns later because of their larger mass and the shielding between the target and detector, as described in Ref. [30]. This separation in arrival time allows beam timing requirements to isolate the prompt signal region while rejecting neutron-induced backgrounds. Although neutron events are therefore not expected to contribute significantly in the physics region of interest, the neutron-dominated sample provides an important data-driven cross-check of the selection.

II. ALP SIGNAL MODEL

This analysis focuses on ALPs that couple predominantly to photons via the effective Lagrangian given in Eq. (1), where $g_{a\gamma}$ denotes the ALP-photon coupling strength, a represents the ALP field, and $F_{\mu\nu}$ ($\tilde{F}^{\mu\nu}$) is the (dual) field strength tensor [37]. The production and detection mechanisms considered here are illustrated in Fig. 1.

$$\mathcal{L}_{a\gamma\gamma} = -\frac{1}{4} g_{a\gamma} a F_{\mu\nu} \tilde{F}^{\mu\nu}. \tag{1}$$

In fixed-target configurations, photon-coupled ALPs are primarily produced via the Primakoff process, in which an incoming photon converts into an ALP through coherent scattering off the nuclear Coulomb field. The corresponding cross section scales approximately as Z^2, providing a significant rate enhancement in high-Z target materials such as the tungsten target at the Lujan Spallation Facility.

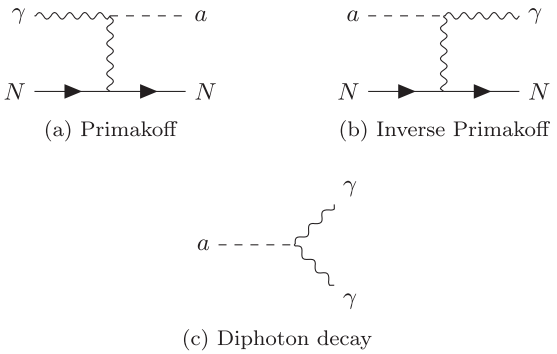


FIG. 1. Diagrams of production and detection processes considered in this analysis. The ALP production mechanism is driven by the Primakoff process, (a), in the target while detection is through a combination of both inverse Primakoff scattering, (b), and diphoton decay, (c), in the detector.

The inverse Primakoff process, where an ALP converts back into a photon within the nuclear electric field, serves as a complementary detection channel. This mechanism is especially critical for low-mass ALPs [$m_a \lesssim \mathcal{O}(10\ \text{keV})$], whose macroscopic decay lengths suppress visible decays within the detector volume.

For heavier ALPs, the dominant experimental signature is the diphoton decay $a \to \gamma\gamma$. In this process, the total decay width is given by Eq. (2), which depends on the coupling strength $g_{a\gamma}$ and the ALP mass m_a. Because the proper lifetime $\tau = 1/\Gamma$ scales as m_a^{-3}, heavier ALPs decay rapidly enough to be observed inside the detector.

$$\Gamma_{a\to\gamma\gamma} = \frac{g_{a\gamma\gamma}^2 m_a^3}{64\pi}. \tag{2}$$

The cross sections for these channels are modeled using the ALPLIB framework [38], building on the previous CCM120 search for ALPs [34]. Similar to that work, the Mark-IV target configuration [39] is fully described in Geant4 [40] to simulate the expected initial gamma-ray flux. For this simulation, particle interactions are modeled with the QGSP_BIC_AllHP physics list, which includes comprehensive treatments of electromagnetic processes, low-energy hadronic interactions, and nuclear deexcitation.

The simulated gamma-ray flux is provided as input to the SIREN injection toolkit [41] to model ALP production, propagation, and detection. The generated final-state photons are subsequently processed through the full detector simulation chain. This includes optical photon production and transport in the Geant4-based detector optical model (presented in Ref. [35]), PMT response simulation, detector timing effects, and data-driven noise overlays. The signal model is constructed on a grid of 100×100 values in mass and coupling, corresponding to 10^4 points in parameter space. The mass range spans 10^{-3} to 10 MeV, while the coupling range extends from 10^{-6} to 10^{-3} GeV^{-1}; both are sampled uniformly in logarithmic scale.

III. EVENT SELECTION LEVERAGING CHERENKOV RADIATION

The primary goal of this work is to incorporate the characteristic temporal and spatial signatures of electromagnetic interactions from ALPs into the event selection procedure, with particular emphasis on Cherenkov radiation. By exploiting both scintillation and Cherenkov light, this analysis demonstrates the potential of a hybrid optical detector to improve signal discrimination in a high-rate beam dump environment.

The event selection begins with the position and energy reconstruction techniques described in the Appendix.

Candidate events must satisfy requirements on reconstructed charge, energy, and fiducial position. Specifically, the reconstructed charge within the first 90 ns must lie between 70 and 3000 photoelectrons (PEs), the reconstructed energy between 0.2 and 10 MeV, the reconstructed radial position within 80 cm of the detector origin, and the reconstructed axial position within $|Z| < 40$ cm. The upper energy threshold is set by the current detector calibration, which is based primarily on an MeV-scale ^{22}Na source [35]. Because the present calibration is limited to MeV-scale data, the reconstructed energy selection is conservatively capped at 10 MeV. Ongoing studies using Michel electrons from cosmic-ray muon decays will extend the validated energy range to approximately 50 MeV.

Although these charge, energy, and fiducial selections substantially reduce the background rate, they do not fully exploit the timing and spatial information contained in the optical signal. Cherenkov photons, despite representing only a small fraction of the detected light, provide complementary information through their prompt arrival times, directional emission, and characteristic spatial distribution.

Applying the Cherenkov-identification techniques developed for calibration data to the full physics dataset is considerably more challenging. For the calibration data, Cherenkov events were identified by requiring hits on uncoated PMTs in the early-time region, which selects the visible component of the Cherenkov emission. Unlike the centrally located sub-MeV ^{22}Na events studied in Ref. [26], ALP interactions span a wider energy range and occur throughout the detector volume, introducing variations in photon propagation, PMT acceptance, and early-time scintillation backgrounds that reduce the purity of the Cherenkov signal.

Reliable identification of the prompt Cherenkov component first requires an accurate determination of the event start time. This analysis uses a constant fraction discrimination (CFD) algorithm, which defines the event time as the point at which the integrated charge reaches 20% of its total value. Using this common timing reference, the discriminating observables described below are evaluated for prebeam steady-state backgrounds, simulated ALP events, and the neutron-dominated control sample to quantify their ability to separate electromagnetic signals from background interactions. For many of these variables, a common time window of -6 to -2 ns relative to the CFD-defined event start time is utilized and will be referred to as the "prompt" time window.

Four discriminating observables are constructed to exploit complementary signatures of electromagnetic interactions. Collectively, they probe (i) enhanced prompt-light collection by the uncoated PMTs, (ii) the directional asymmetry of the earliest detected photons, (iii) pulse-shape differences arising from prompt and delayed optical components, and (iv) the overall spatial extent of the reconstructed event. Together, they provide additional discrimination between electromagnetic signal events and background interactions, highlighting the advantages of hybrid optical detection in background rejection. In the rest of this section, these variables are evaluated on the prebeam steady-state background sample, the entire ALP Monte Carlo sample (10^{-3} MeV $< m_a <$ 10 MeV), and the neutron-dominated sample.

A. Number of hits on the uncoated PMTs

The first variable used for event discrimination is the number of reconstructed PE hits recorded by the uncoated PMTs within the prompt time window. This time interval is selected to maximize sensitivity to the prompt optical signal, where the contribution from Cherenkov photons is expected to be greatest.

The uncoated PMTs are particularly well suited for identifying Cherenkov light during this early-time interval. Because they lack a wavelength-shifting coating, they have little sensitivity to the vacuum ultraviolet scintillation photons and instead primarily detect wavelength-shifted scintillation light together with the visible component of Cherenkov radiation. Consequently, prompt hits on the uncoated PMTs are preferentially associated with visible Cherenkov light, whereas scintillation photons are delayed by both the intrinsic scintillation emission time constant and the additional time required for the photons to propagate to the TPB, undergo wavelength shifting, and reach the PMTs.

Figure 2 compares the distributions of the prompt uncoated PMT hit multiplicity for prebeam data, simulated ALP events, and neutron-dominated data. Both the prebeam and ALP samples are markedly peaked at zero hits, indicating that relatively few photons are detected in this narrow prompt window. Nevertheless, the ALP simulation contains a noticeably larger population of events with one or more hits, as expected from the prompt Cherenkov light associated with electromagnetic final states.

The neutron-dominated sample exhibits clearly different behavior, producing a much wider distribution with a

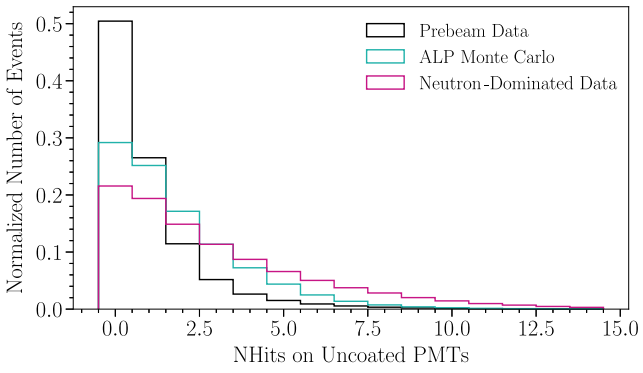


FIG. 2. Number of hits on the uncoated PMTs for the prebeam data events, ALP Monte Carlo, and neutron-dominated data events.

substantially larger number of events containing multiple prompt hits on the uncoated PMTs. This trend is consistent with the increased scintillation light generated by higher-energy interactions, which produces greater spillover into the early-time region.

These results demonstrate that the simple hit-count threshold on the uncoated PMTs employed in the ^{22}Na calibration analysis (the subject of Ref. [26]) provides limited discrimination for the ALP search. At higher energies and for interactions occurring throughout the detector volume, the prompt uncoated PMT hit multiplicity alone provides only limited separation between Cherenkov-rich electromagnetic events and scintillation-dominated backgrounds. This limitation motivates combining multiple complementary observables within a multivariate discrimination framework.

B. Spread in the directionality

The next variable utilizes the spatial anisotropy of Cherenkov radiation for signal isolation. Unlike scintillation emission which is isotropic, Cherenkov light is directional, offering a robust metric for event discrimination. To calculate this, we restrict the analysis to hits within the prompt time window. Within this interval, we compute the directionality metric C, defined in Eq. (3), which evaluates the angular uniformity of the triggered PMTs.

Specifically, C represents the squared magnitude of the charge-weighted average of unit vectors pointing from the reconstructed interaction vertex ($\vec{x}_{\text{vertex}}$) to each hit PMT position ($\vec{x}_i$) where q_i is the charge of the ith PMT and $Q = \sum_i q_i$ represents the cumulative charge inside the prompt window:

$$C = \left|\left| \sum_i \left(\frac{q_i}{Q}\right) \frac{\vec{x}_i - \vec{x}_{\text{vertex}}}{||\vec{x}_i - \vec{x}_{\text{vertex}}||} \right|\right|^2. \tag{3}$$

Bounded between 0 and 1, C approaches 0 for uniform, isotropic distributions and nears 1 for highly collinear light emission. Electromagnetic events containing Cherenkov light are therefore expected to exhibit larger values of C than isotropic scintillation-dominated backgrounds.

Figure 3 contrasts the C distributions for prebeam data, ALP simulations, and neutron-dominated events. The ALP signal is shifted toward larger values (centered around $C \sim 0.3$), reflecting the directional Cherenkov light produced by electromagnetic final states. In contrast, neutron-dominated events cluster at smaller values (around $C \sim 0.15$), where the optical signal is primarily composed of isotropic scintillation light. The prebeam sample occupies an intermediate region (around $C \sim 0.2$), consistent with a mixture of electromagnetic backgrounds, detector noise, and other background processes. This variable isolates a fundamental physical difference between signal and background, making it a key component of the joint likelihood discriminant.

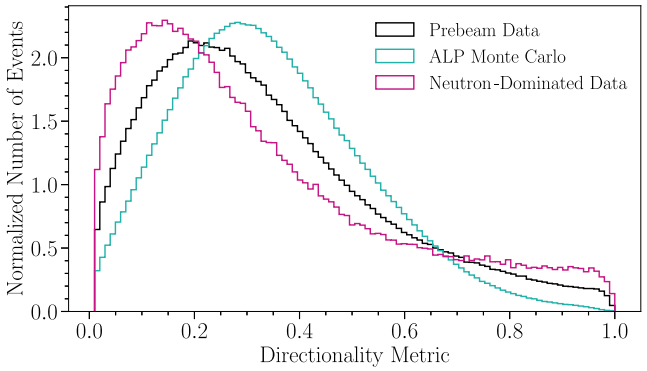


FIG. 3. Distribution of the directionality metric C, defined in Eq. (3), across the analyzed datasets. The ALP Monte Carlo events show higher values due to directional photon emission, while neutron-dominated events skew lower, as expected for isotropic light.

C. Pulse shape ratio

The third variable included in the likelihood ratio is based on the temporal structure of the detected optical pulse. In traditional liquid argon scintillation detectors, pulse shape discrimination (PSD) is commonly used to distinguish nuclear recoils from electromagnetic interactions by taking advantage of differences in their scintillation decay profiles, which arise from variations in ionization density (dE/dx) [42–45]. In the present analysis, however, a conventional PSD technique cannot be implemented because of the operating conditions of the experiment.

The CCM experiment operates in a pulsed beam environment with substantial beam-related activity. Conventional PSD integrates charge over several microseconds to capture the long-lived triplet-state scintillation emission, but such long integration windows would be highly susceptible to pileup. The charge integration window is therefore limited to $\mathcal{O}(100\ \text{ns})$, balancing photostatistics against pileup contamination.

Although this integration window precludes conventional PSD, the pulse time evolution still provides useful discriminating information. A pulse shape ratio (PSR) is constructed as the ratio of the charge collected during the early portion of the pulse to the total prompt charge, as defined in Eq. (4). All timing quantities are measured relative to the CFD-determined event start time.

$$\text{PSR} = \frac{Q(0\ \text{ns} < t < 20\ \text{ns})}{Q(0\ \text{ns} < t < 90\ \text{ns})}. \tag{4}$$

The PSR distributions for simulated ALP events, prebeam data, and neutron-dominated events are shown in Fig. 4. The simulated ALP sample exhibits a narrow peak near ~0.6, while the prebeam and neutron-dominated distributions are broader and shifted toward smaller values, consistent with a larger delayed-light contribution. The

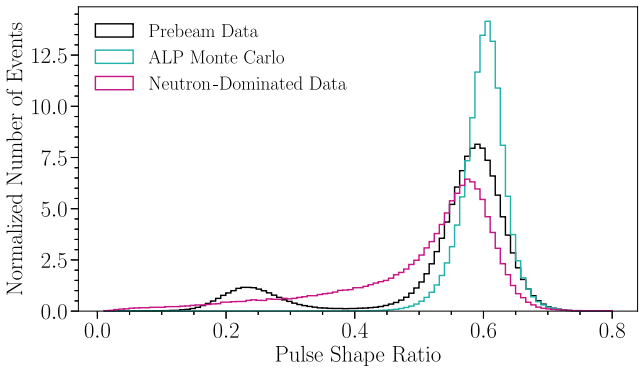


FIG. 4. Pulse shape ratio distributions across the ALP, prebeam, and neutron-dominated datasets. While the ALP Monte Carlo events are narrowly peaked around 0.6, the prebeam and neutron-dominated datasets are broader and skewed toward lower values.

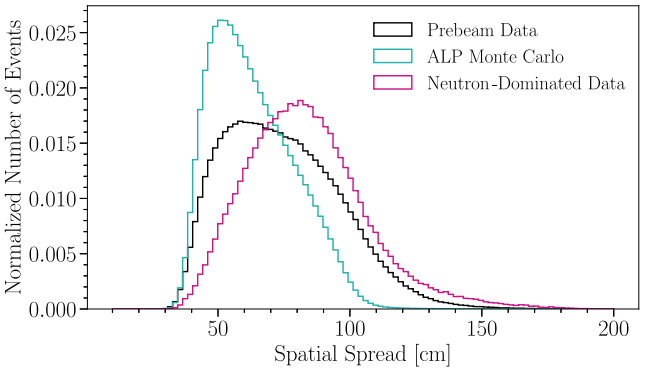


FIG. 5. Spatial extent of the early reconstructed charge distribution. The rms, defined in Eq. (5), distinguishes localized electromagnetic energy depositions from more spatially extended backgrounds.

prebeam data also contain a distinct population near ∼0.25, attributed to electronics noise and overlapping events. Although the PSR is not a conventional PSD variable, it provides complementary timing information that improves signal to background discrimination.

D. Spatial spread of events

The final input to the likelihood ratio is designed to quantify the spatial distribution of the detected charge. This variable takes advantage of the differing event topologies expected for signal and background interactions. ALP-induced events, which produce one or two final-state gamma rays, generally deposit their energy in relatively compact electromagnetic showers. By comparison, neutron backgrounds often undergo multiple scattering and secondary interactions within the detector, resulting in a more diffuse spatial distribution of charge.

$$\mathrm{rms} = \sqrt{\frac{\sum_i q_i ||\vec{x}_i - \vec{x}_{\mathrm{vertex}}||^2}{\sum_i q_i}}. \tag{5}$$

This behavior is quantified using the charge-weighted root-mean-square (rms) distance of the PMT hits from the reconstructed interaction vertex, as defined in Eq. (5). Here, $\vec{x}_i$ denotes the position of the ith PMT, $\vec{x}_{\mathrm{vertex}}$ is the reconstructed event vertex, and q_i is the charge measured by that PMT.

The rms is calculated using only hits within the prompt time window. Restricting the calculation to this interval suppresses contributions from delayed and wavelength-shifted scintillation light, which would otherwise broaden the measured charge distribution. As a result, the rms provides a measure of the characteristic spatial scale of the prompt light: compact energy depositions produce smaller rms values, whereas extended or multisite interactions yield larger values.

The resulting rms distributions for the ALP Monte Carlo, prebeam data, and neutron-dominated data are shown in Fig. 5. The simulated ALP events are concentrated at lower rms values, with a peak near ∼50 cm, reflecting the compact topology of electromagnetic interactions. In contrast, the neutron-dominated sample peaks around ∼80 cm and exhibits a broader distribution, consistent with the larger spatial extent expected from multiple neutron scatters. The prebeam data fall between these extremes, with a maximum near ∼60 cm and an extended tail toward larger rms values, indicative of a mixture of physical backgrounds and nonphysical events.

E. Likelihood ratio test

The four observables introduced above are combined into a single discriminant using a log-likelihood ratio (LLR). For each variable, x_k, probability density functions (PDFs) are constructed under both the signal and background hypotheses. Signal PDFs are obtained from the simulated ALP sample, while background PDFs are derived from the prebeam data.

$$\mathrm{LLR} = \sum_k \log(p_{\mathrm{sig}}(x_k)) - \log(p_{\mathrm{bkg}}(x_k)). \tag{6}$$

The likelihood for a given event is computed as the product of the individual PDFs for the four observables. The LLR, given by Eq. (6), is therefore equal to the logarithm of the ratio between the signal and background likelihoods, where $p_{\mathrm{sig}}(x_k)$ and $p_{\mathrm{bkg}}(x_k)$ denote the corresponding signal and background PDFs evaluated at the measured value of the kth observable.

The PDFs are generated from normalized histograms of each variable and interpolated to obtain continuous probability estimates. To prevent numerical instabilities in sparsely populated regions, a small regularization term is included in each PDF.

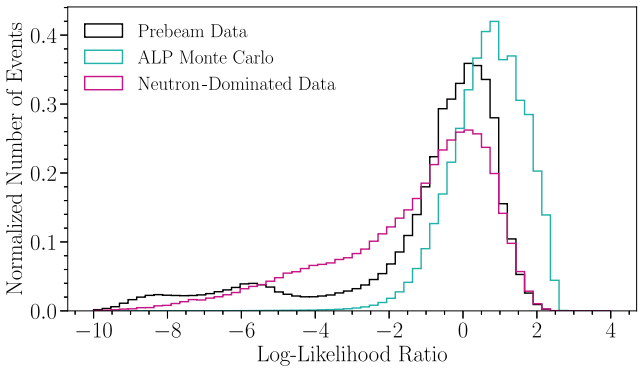


FIG. 6. LLR discriminant formed from the four event-level observables described previously. Signal events preferentially occupy larger LLR values than either the prebeam or neutron-dominated background samples. The analysis defines the signal region by requiring LLR > 1, balancing signal efficiency against background rejection.

Figure 6 presents the resulting LLR distributions for the ALP simulation, prebeam data, and neutron-dominated data. Signal events preferentially populate larger LLR values, with a broad maximum near 1, whereas both background samples are concentrated at lower values with more pronounced tails. A requirement of LLR > 1 is adopted to define the final signal region, as it provides good background rejection while retaining high signal acceptance.

F. Final event selection

The event selection is optimized to reject prebeam backgrounds while maintaining sensitivity to the low-energy electromagnetic signatures expected from ALP decays. The complete selection and the corresponding efficiencies for the prebeam sample are summarized in Table I.

Basic charge, reconstructed energy, and reconstructed position cuts are applied, as described at the beginning of Sec. III. After those cuts, the final selection criterion is imposed on the log-likelihood ratio described in the previous section. Events with LLR > 1 are retained, providing additional separation between signal-like electromagnetic interactions and residual backgrounds. Together, these requirements reduce the prebeam sample to 0.49% of its original size. The previous iteration of this analysis using the prototype detector, described in Ref. [34], achieved 3.2% efficiency for the prebeam data sample. The detector upgrades allowing use of the Cherenkov signal enhanced the steady-state background rejection capabilities by approximately sixfold.

TABLE I. Selection efficiencies of cuts utilized in this analysis for the prebeam data sample.

Cut	Efficiency [%]
Charge	49.75
Position	18.29
Energy	80.93
Charge + position + energy	7.19
Charge + position + energy + LLR	0.49

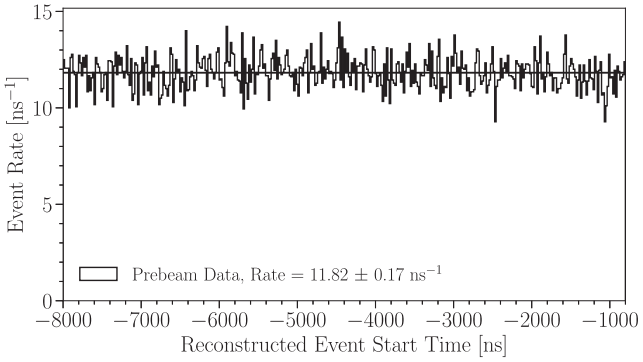


FIG. 7. Time distribution of prebeam events after all selection criteria are applied. A fit to a constant rate yields an expected background of 11.82 ± 0.17 events per ns.

Figure 7 shows the time distribution of the surviving prebeam events after all selections have been applied. The prebeam time window is defined as reconstructed start times between -8000 and -800 ns relative to an external beam current monitor. The event rate, integrated over the full exposure of 1.23×10^{21} POT, is consistent with a time-independent background.

A fit with a constant function yields a background rate of 11.82 ± 0.17 events/ns. The fit quality is evaluated using Pearson's χ^2 test, giving $\chi^2 = 427.76$ for 399 degrees of freedom, corresponding to a reduced χ^2 of 1.07. The result is consistent with a uniform background rate and shows no statistically significant residual time dependence following the event selection.

The reconstructed energy spectrum of the selected prebeam sample is presented in Fig. 8. Most events cluster

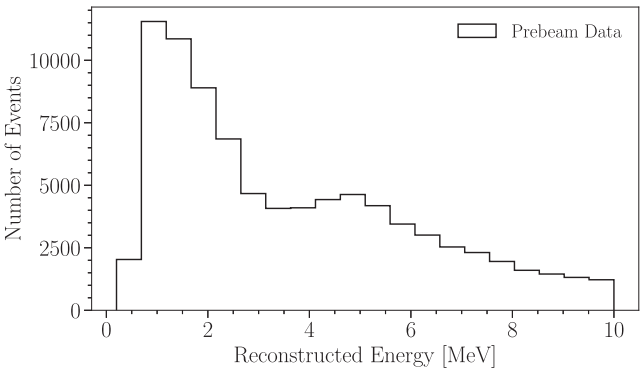


FIG. 8. Reconstructed energy spectrum of prebeam events surviving the full event selection. The feature near 5 MeV may be associated with gamma rays from neutron capture on argon, although additional study is required to establish its origin.

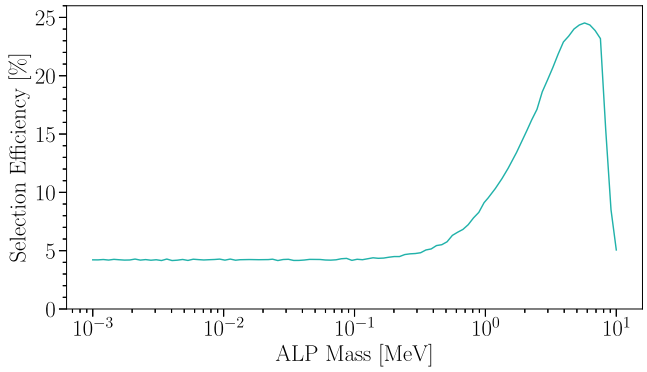


FIG. 9. Overall signal selection efficiency as a function of ALP mass. The efficiency increases with mass, reaching approximately 25% near 6 MeV before falling rapidly because of the upper reconstructed energy requirement of 10 MeV.

near 1 MeV, consistent with low-energy electromagnetic backgrounds that satisfy the Cherenkov-enhanced event selection.

A second feature is visible near 5 MeV. One possible explanation is gamma-ray emission following neutron capture on argon through the reaction $^{40}\mathrm{Ar}(n,\gamma)^{41}\mathrm{Ar}$, which produces deexcitation gamma rays with energies totaling up to approximately 6.1 MeV. These neutrons could originate from previous beam spills and remain in the experimental hall before thermalizing and capturing in the detector. Although this hypothesis is compatible with the observed energy scale, dedicated studies are needed to determine the origin of this feature.

The overall signal efficiency after all selection requirements depends strongly on the ALP mass and is shown in Fig. 9. For ALP masses below approximately 0.3 MeV, the efficiency remains nearly constant at about 4%, reflecting the small visible energy deposited by these events. As the ALP mass increases, the larger electromagnetic energy deposition improves reconstruction and event selection, causing the efficiency to rise steadily. The efficiency reaches a maximum of roughly 25% near 6 MeV before decreasing sharply at higher masses, where an increasing fraction of signal events fail the upper reconstructed energy requirement of 10 MeV.

IV. AXIONLIKE PARTICLE FIT

The fit is performed in the two-dimensional space of reconstructed energy and event start time. Including timing information enhances the separation of beam-correlated ALP signals from the uniform steady-state background. A frequentist scan is performed over a grid of 10^4 ALP parameter points, in addition to the background-only hypothesis. At each point, simulated ALP events are propagated through the full detector simulation, reconstruction, and event selection to produce signal templates. These templates are compared to the observed data by minimizing an effective binned likelihood that accounts for finite Monte Carlo statistics [46]. Confidence intervals are derived using Wilks' theorem [47].

A. Fitting procedure

The likelihood fit is performed using two-dimensional templates in reconstructed energy and event start time. The energy axis is divided into 12 uniform bins spanning 0.2–10 MeV, while the time axis consists of 20 uniform bins covering −600 to −424 ns relative to an external beam current monitor. This time period defines the physics region of interest (ROI), beginning with the earliest possible arrival of relativistic particles at the detector and ending with the onset of neutron backgrounds. The start of the ROI is determined using an external EJ-301 scintillator detector to monitor beam-related activity in the experimental hall. By identifying the statistically significant increase in activity, the arrival time of prompt relativistic particles is established. Pulse-level activity distributions in the CCM200 detector are then used to identify the onset of neutron backgrounds, which defines the end of the ROI. For further details, see Ref. [30].

Observed data events in the time ROI and the energy range of the analysis are filled directly into the two-dimensional (E,t) histogram, demonstrated in Fig. 10. To reduce statistical fluctuations in the signal and background templates, the energy distributions are obtained from simulated ALP or prebeam events, while the time distributions are modeled analytically and normalized to the corresponding energy spectra. This hybrid approach minimizes the impact of finite Monte Carlo statistics on the time dimension.

The selected prebeam sample is consistent with a constant event rate in time. Consequently, the expected background time distribution is modeled with a uniform PDF.

The expected signal time distribution is constructed by convolving the simulated event time offsets, which account

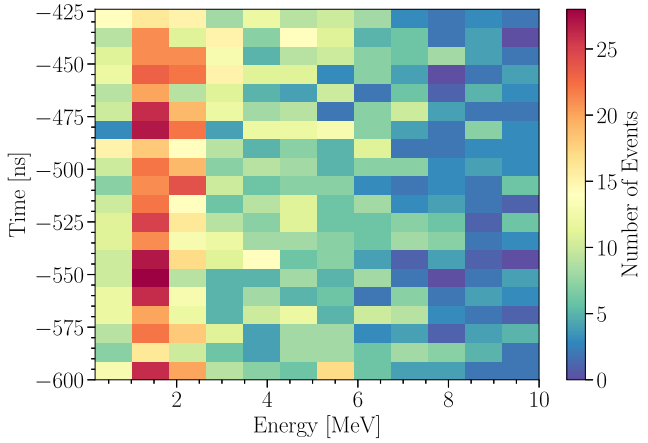


FIG. 10. Two-dimensional reconstructed energy and event start-time distribution of the observed data.

for the ALP time of flight to the detector along with reconstruction-level timing offsets, with the time profile of the beam pulse. The beam pulse has a characteristic triangular time profile that is 290 ns wide at the base, shown in Ref. [36]. For each simulated event, the probability of falling within a given time bin is computed by integrating the beam time profile over the bin boundaries, and these probabilities are summed using the corresponding event weights.

B. Sources of systematic uncertainties

As a proof-of-concept analysis, the available statistics are limited, and the measurement is therefore dominated by statistical uncertainties. Nevertheless, several sources of systematic uncertainty are considered and incorporated into the likelihood through nuisance parameters that are profiled during the fit. The dominant contributions arise from the steady-state background normalization, the total accumulated POT, and the absolute beam timing.

1. Background normalization

The selected prebeam sample exhibits a time-independent event rate (Fig. 7), with a best-fit value of 11.82 ± 0.17 events/ns. The background normalization is therefore treated as a nuisance parameter with a Gaussian prior centered at 11.82 events/ns and a standard deviation of 0.17 events/ns. This allows the fit to vary the overall background level within the measured uncertainty while constraining deviations from the nominal expectation.

2. Total protons on target

The expected ALP signal scales linearly with the total POT accumulated during data taking. The POT is determined from a calibration between an external beam current monitor and the beam current reported by the LANSCE accelerator division, resulting in an overall normalization uncertainty of approximately 5%. This uncertainty is incorporated through a multiplicative signal normalization parameter with a Gaussian prior centered at unity and a standard deviation of 0.05.

3. Beam timing

The time distribution of the signal prediction depends on the relative timing between the proton beam and the detector. The earliest arrival time of prompt beam-related particles is known to within approximately 30 ns, due to statistics of the ROI start time determination, producing a corresponding uncertainty in the predicted ALP timing distribution. To account for this effect, the beam start time is allowed to vary during the fit with a Gaussian prior centered at -600 ns and a width of 30 ns.

C. Fit results

The fit is performed by minimizing the negative log-likelihood while profiling the three nuisance parameters introduced in the previous section. The effective likelihood, including Gaussian pull terms for these nuisance parameters, is given in Eq. (7). This formalism depends on the observed counts k and the weighted Monte Carlo quantities $\alpha = \frac{(\sum_i w_i)^2}{\sum_i w_i^2} + 1$ and $\beta = \frac{\sum_i w_i}{\sum_i w_i^2}$, where w_i is the weight assigned to each simulated event [46]. The pull terms correspond to Gaussian priors centered at μ_i with standard deviations σ_i for each nuisance parameter x_i.

$$\mathcal{L}_{\text{eff}}(k,\alpha,\beta) = \prod_{(E,t)} \frac{\Gamma(k+\alpha)}{k!\Gamma(\alpha)} \frac{\beta^\alpha}{(1+\beta)^{k+\alpha}} \times \prod_i \exp\left(-\frac{(x_i-\mu_i)^2}{2\sigma_i^2}\right). \tag{7}$$

Under the background-only hypothesis, the fit returns a negative log-likelihood of 578.28 and a profiled background rate of 11.74 events/ns, consistent with the measured value within its 1σ uncertainty range.

Allowing for an ALP signal, the best-fit point occurs at $m_a = 0.18$ MeV and $g_{a\gamma} = 1.63 \times 10^{-4}$ GeV^{-1}, with a negative log-likelihood of 577.34. The profiled nuisance parameters remain close to their nominal values, yielding a POT scaling of 0.99, a beam start time of -584.52 ns, and a background rate of 11.67 events/ns. The small shifts in these parameters indicate that the data are well described without requiring substantial deviations from the nominal model.

The best-fit signal yields a decrease in the negative log-likelihood of 0.94 relative to the background-only hypothesis. Under Wilks' theorem, for two d.o.f. corresponding to the ALP mass and coupling, this corresponds to a local significance of 0.86σ. The observed improvement is therefore compatible with a statistical fluctuation, and no evidence for ALP production is found.

Figure 11 shows the reconstructed energy projection of the fit along with the observed data in black. The background-only prediction is shown in cyan, while the best-fit signal-plus-background model is shown in magenta, with the dashed curve indicating the signal component alone. Note that the signal-only component (dashed magenta) is magnified by a factor of 10 to show the features more clearly. The preferred signal produces a modest excess between approximately 1 and 3 MeV, but the statistical uncertainty in this region is sufficiently large that the improvement in fit quality is minimal.

The corresponding event start-time projection is shown in Fig. 12. The background expectation is uniform across the analysis window, whereas the signal model exhibits the characteristic triangular timing profile produced by the proton beam pulse, beginning near -584 ns and extending

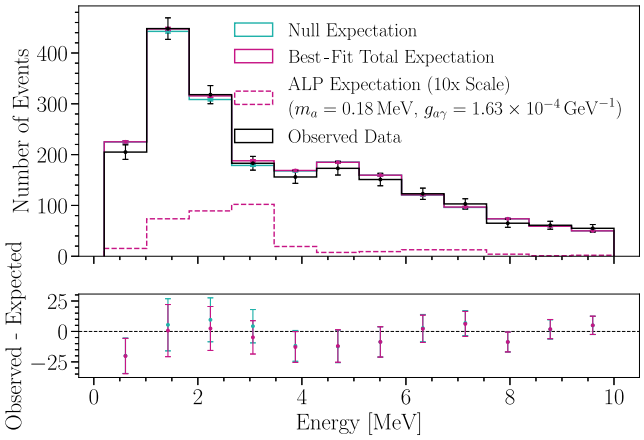


FIG. 11. Projected reconstructed energy distribution for the observed data (black), the background-only fit (cyan), and the best-fit signal-plus-background model (magenta). The best fit ALP signal is at $m_a = 0.18$ MeV and $g_{a\gamma} = 1.63 \times 10^{-4}$ GeV^{-1}. The dashed magenta line shows the ALP expectation at 10× scale for visibility. The lower panel shows the residuals relative to the observed data.

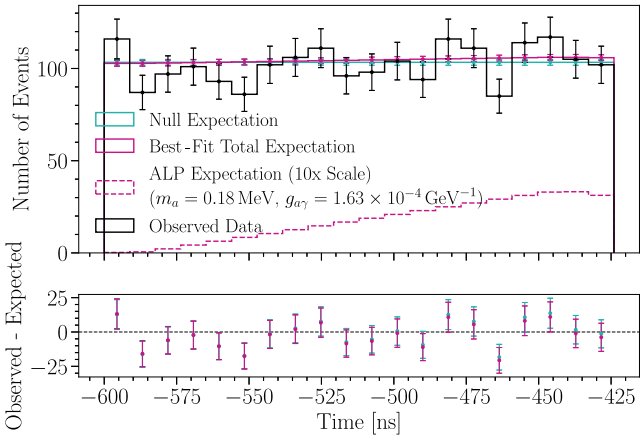


FIG. 12. Projected event start-time distribution for the observed data (black), the background-only fit (cyan), and the best-fit signal-plus-background model (magenta). The best fit ALP signal is at $m_a = 0.18$ MeV and $g_{a\gamma} = 1.63 \times 10^{-4}$ GeV^{-1}. The dashed magenta line shows the ALP expectation at 10× scale for visibility. The lower panel shows the residuals relative to the observed data.

over approximately 290 ns. The observed timing distribution shows no statistically significant preference for this signal-like structure.

Finally, Fig. 13 presents the observed 90% confidence level exclusion limits together with the expected sensitivity, where the 90% confidence level corresponds to $\Delta\text{NLLH} = 2.30$. In addition to the previous CCM120 search for ALPs, the allowed parameter space for the Kim-Shifman-Vainshtein-Zakharov QCD axion model is denoted by the dashed black lines [4,5]. Existing constraints on the ALP parameter space are also shown for

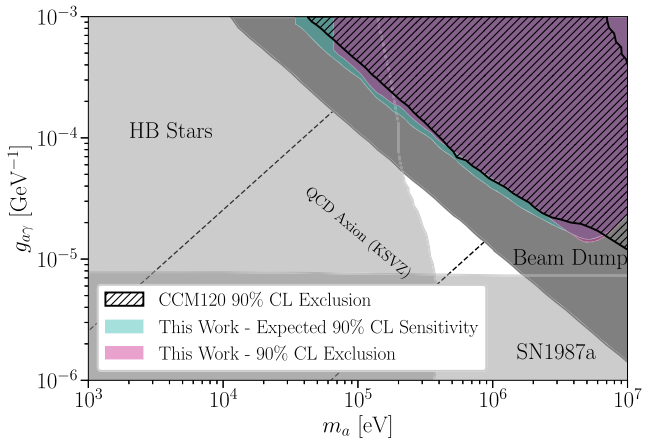


FIG. 13. Observed (magenta) and expected (cyan) 90% confidence level exclusion limits on the ALP parameter space obtained in this analysis. No statistically significant evidence for an ALP signal is observed. Despite using only approximately 70% of the proton exposure of the previous CCM120 search, this analysis extends the excluded parameter space through improved background rejection enabled by Cherenkov-based event discrimination.

comparison. Beam dump experiments provide the strongest laboratory limits in this region by probing similar ALP production and decay processes at comparable masses and couplings [48,49]. Astrophysical constraints are also considered, including bounds from energy loss arguments in horizontal branch stars and observations of the neutrino signal from SN1987a [50–57]. These limits provide strong complementary coverage, especially at low masses and couplings, but depend on assumptions about stellar models and ALP interactions in astrophysical environments.

Although this analysis uses approximately 70% of the integrated POT compared to the previous CCM120 search, it excludes additional regions of ALP parameter space due to increased suppression of steady-state backgrounds. This demonstrates the improved sensitivity achieved through the hybrid Cherenkov-scintillation detection and its enhanced background rejection.

V. SUMMARY AND FUTURE PROSPECTS

This analysis presents the first search for axionlike particles using a 10-ton liquid argon hybrid Cherenkov-scintillation optical detector in a beam dump environment. By exploiting four complementary observables based on the distinct timing, wavelength response, angular structure, and topology of Cherenkov and scintillation light, this work demonstrates the ability of hybrid optical detection to enhance discrimination between electromagnetic signal events and beam-related backgrounds.

Table II compares the detector configurations and analysis performance of the CCM120 ALP search [34] and this work. Despite using a smaller exposure, this work achieves a substantially lower prebeam background

TABLE II. Comparison between CCM120 and CCM200 (this work). The prebeam efficiency defines the rejection of steady-state backgrounds, discussed in Sec. III. The efficiency for $m_a = 1$ MeV ALPs between the two analyses is provided for comparison purposes at one point in the parameter space.

Detector configuration	Hybrid detection	Photocathode coverage (%)	Exposure	Prebeam eff. (%)	1 MeV ALP eff. (%)
CCM120	No	30	1.79×10^{21} POT	3.2	19.0
CCM200	Yes	50	1.23×10^{21} POT	0.49	9.22

efficiency, with a factor of ∼6 improvement compared to CCM120. This improvement is enabled by the hybrid Cherenkov-scintillation detection capabilities of CCM200, which provide additional event discrimination beyond scintillation information alone. The improved background rejection compensates for the reduced exposure and signal efficiency, resulting in comparable or improved ALP sensitivity across the presented mass range.

The sensitivity of the current search is optimized for low-energy signals, reflecting the detector calibration based on the ^{22}Na source presented in Ref. [35]. Ongoing calibration efforts using Michel electrons will extend the validated energy range to approximately 50 MeV, enabling future searches to probe higher-mass ALPs with improved sensitivity. The collaboration is also working on direction reconstruction that will be especially useful to isolate boosted signals expected from higher energy ALPs. Additionally, future analyses will utilize the full CCM200 physics dataset, comprising approximately 3×10^{21} POT.

Beyond this specific search, this work demonstrates that hybrid optical detectors can substantially improve sensitivity for future accelerator-based searches for weakly interacting particles.

ACKNOWLEDGMENTS

We acknowledge the support of the Los Alamos National Laboratory Laboratory Directed Research and Development and the U.S. Department of Energy Office of Science funding. We also acknowledge support from the Los Alamos Neutron Science Center Lujan Center and Los Alamos National Laboratory's Accelerator Operations and Technology (AOT) division. This research used resources provided by the Los Alamos National Laboratory Institutional Computing Program, which is supported by the U.S. Department of Energy National Nuclear Security Administration under Contract No. 89233218CNA000001. D. A. N. is supported by the U.S. National Science Foundation Graduate Research Fellowship under Grant No. 2141064 and MIT School of Science. A. A. A.-A., C. F. M.-A., J. C. D., and M. C.-E. acknowledge support from Dirección General de Asuntos del Personal Académico-Universidad Nacional Autónoma de México Grant No. PAPIIT-IN104026. We also thank the SubMIT supercomputer and physics analysis facility at Massachusetts Institute of Technology [58] for the generous use of their computing resources.

DATA AVAILABILITY

There are no publicly available research data or software supporting this manuscript. Requests for further information or data should be sent to the authors.

APPENDIX: POSITION AND ENERGY RECONSTRUCTION

Position reconstruction is performed using the GraphNeT machine learning framework [59], which was adapted for the CCM200 detector geometry and achieves a spatial resolution of approximately 5 cm in each dimension. The reconstructed position is then used together with the measured charge deposition to estimate the event energy, yielding an energy resolution of approximately 12% at 1 MeV.

1. Position reconstruction

Position reconstruction provides strong discrimination between signal and background events. Beam-related and cosmic-ray backgrounds predominantly originate outside the detector and therefore reconstruct near its boundaries, while signal events are expected to be uniformly distributed throughout the active volume. Improving the vertex resolution directly enhances the signal-to-background ratio.

To improve upon the charge-weighted PMT position-based reconstruction used in previous CCM120 analyses, a graph neural network-based position reconstruction algorithm was developed using the open-source GraphNeT framework [59]. The model employs a DeepIce transformer backbone [60], which naturally handles the irregular geometry and sparse optical readout of the detector.

The network is trained on reconstructed photoelectron pulses, using the pulse time, charge, PMT position, and PMT coating status as input features. Embedded pulse features are processed by transformer layers with self-attention, allowing the network to learn the spatial and temporal correlations of detected light. Event-level representations are formed using a learned classification token, whose output is passed to a regression head that predicts the interaction vertex.

Training is performed with the LogCosh loss function, detailed in Eq. (A1). In this loss function, $\Delta x_i = \hat{y}_i - y_i$, for y_i representing the true position and $\hat{y}_i$ is the model prediction. This loss was optimized using the Adam optimizer [61].

$$\mathrm{LogCosh}(\Delta x_i) = \sum_{i=1}^{n} [\Delta x_i + \ln(1 + e^{-2\Delta x_i}) - \ln(2)]. \tag{A1}$$

The position reconstruction model is trained on approximately 6×10^6 simulated electron events generated uniformly throughout the detector volume, including the veto region, with isotropic directions and energies uniformly distributed between 0.1 and 30 MeV. The training target is the true electron injection position. To accurately reproduce detector response, the simulation incorporates the full data acquisition chain, including measured board-to-board timing offsets, PMT electron transit time calibrations, and 2 ns waveform digitization. A data-driven noise model is then overlaid by sampling prebeam pulse series from physics data, reproducing the steady-state PMT noise observed during data taking.

The simulated events are processed using the same event reconstruction chain as the physics data to avoid biasing the selection. Additionally, the reconstructed pulse series are rebinned to 6 ns time bins to reduce data quantity. To improve robustness against detector systematics, each training event includes random variations corresponding to the measured 10% uncertainty in the scintillation light yield [35] and a small ($\sigma = 0.5$ ns) Gaussian time jitter applied to individual pulse times, reducing sensitivity to digitization artifacts.

a. Position reconstruction performance

The model is evaluated on a held-out validation sample comprising 10% of the simulated events. Figure 14 compares the reconstructed and true interaction positions in each spatial dimension after applying the standard fiducial selection ($r < 80$ cm and $|z| < 40$ cm). The reconstructed positions exhibit a strong linear correlation with the true positions, with no significant bias observed across the fiducial volume.

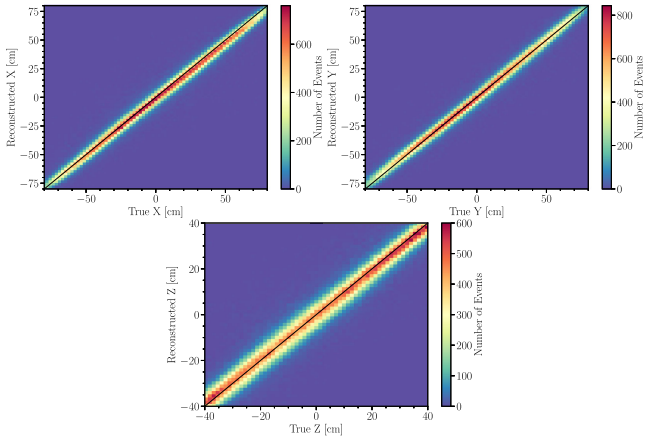


FIG. 14. Two-dimensional distributions comparing the true and reconstructed event positions for the GraphNeT position reconstruction model. The x, y, and z coordinates are shown in the top left, top right, and bottom panels, respectively. The diagonal line denotes the ideal reconstruction.

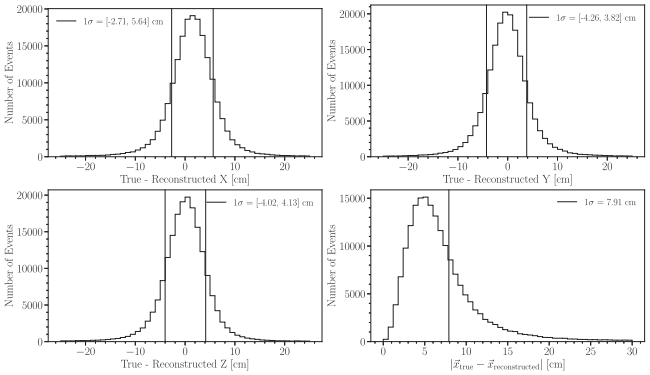


FIG. 15. Residual distributions for the reconstructed event position. The x, y, and z coordinates each exhibit a resolution of approximately 5 cm at the 1σ level, while the combined three-dimensional position resolution is 7.91 cm.

The corresponding residual distributions, shown in Fig. 15, are used to quantify the spatial resolution. The reconstruction achieves an approximately isotropic resolution of ∼5 cm in each coordinate, corresponding to a three-dimensional position resolution of 7.9 cm at the 1σ level. This represents roughly a fourfold improvement over the charge-weighted reconstruction previously used in CCM120 analyses, enabling more precise fiducialization and improved rejection of backgrounds originating near the detector boundaries.

The position reconstruction is validated using ^{22}Na calibration data collected with the source deployed at the detector center. Figure 16 compares the reconstructed position distributions for source data and a background-only dataset acquired with the source removed. As expected, the source data exhibit a pronounced peak near

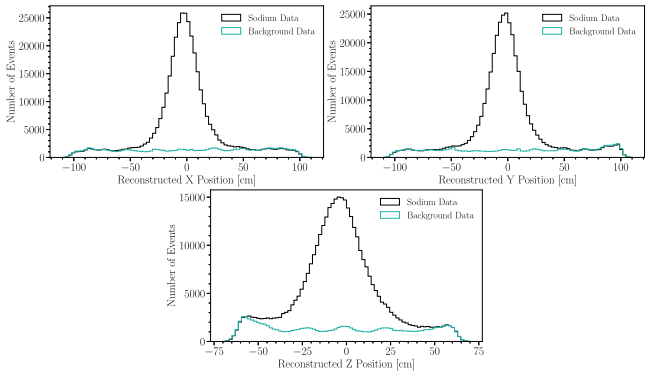


FIG. 16. Reconstructed event positions for the ^{22}Na calibration and background datasets. The calibration sample (black) is concentrated near the detector origin, as expected, whereas the background events (cyan) are distributed approximately uniformly within the active volume.

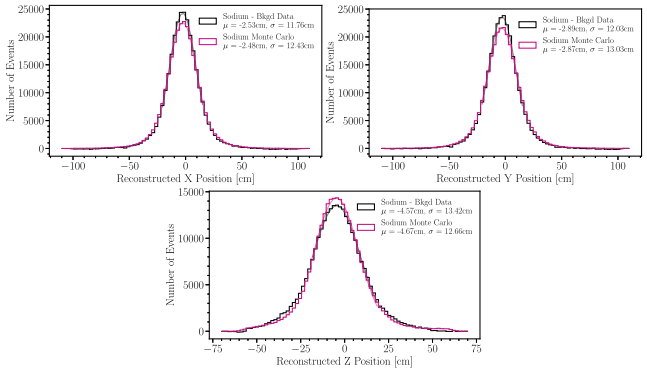


FIG. 17. Reconstructed position distributions for background-subtracted ^{22}Na calibration data and Monte Carlo simulation. Gaussian fits are used to extract the reconstructed source position and width, with agreement between data and simulation at the $\mathcal{O}(1\ \mathrm{cm})$ level in all dimensions.

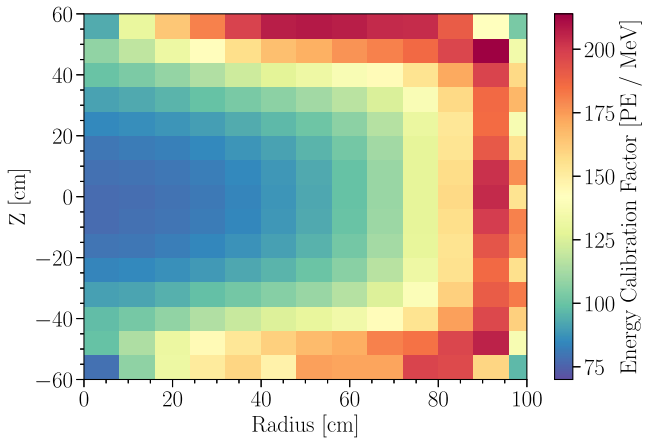


FIG. 18. Energy calibration factor as a function of reconstructed radius and Z position. The detector response varies from approximately 80 PE/MeV at the center to 200 PE/MeV near the edges, motivating the use of a position-dependent energy reconstruction.

the detector origin, while the background sample is approximately uniform throughout the detector volume.

To isolate the ^{22}Na contribution, the background-only sample is normalized to the same exposure and subtracted from the source data. The resulting distributions are compared with Monte Carlo simulation in Fig. 17. Gaussian fits to the reconstructed positions show agreement between data and simulation at the $\mathcal{O}(1\ \mathrm{cm})$ level in all three spatial dimensions. The observed widths, of order 10 cm, are larger than the expected position resolution obtained from simulated electron events because the ^{22}Na source emits gamma rays, which travel finite distances before producing secondary electrons. The agreement between data and simulation demonstrates that the reconstruction accurately reproduces the source position and detector response in calibration data.

2. Energy reconstruction

In addition to position, energy reconstruction is also essential for event characterization. The deposited energy is inferred from the observed scintillation light, with the total number of detected PEs serving as a proxy for the true energy deposition. However, the light yield depends on the interaction position within the detector because of variations in light collection efficiency, optical propagation, and PMT response.

To account for these effects, the energy reconstruction incorporates the reconstructed interaction position together with the observed charge to estimate the deposited energy. This position-dependent approach significantly improves upon the uniform charge-to-energy calibration used in previous CCM120 analyses.

The deposited energy is estimated from the charge collected within the first 90 ns of an event using a position-dependent calibration. A large sample of simulated electrons, uniformly distributed throughout the detector with energies between 0.1 and 30 MeV, is used to derive this calibration. After reconstructing each interaction vertex using the techniques described previously, the detector is divided into 8 cm voxels, and the median charge-to-energy ratio (Q/E) is computed in each voxel. A spline interpolation is then used to obtain a continuous energy calibration as a function of reconstructed position.

The resulting calibration, shown in Fig. 18, varies smoothly across the detector, increasing from approximately 80 PE/MeV at the center to 200 PE/MeV near the detector boundaries because of improved light collection. This variation of more than a factor of 2 highlights the importance of accounting for position when reconstructing the deposited energy. Restricting the charge integration to the first 90 ns also reduces contamination from delayed light and pileup, providing a robust energy estimate for the low-energy, contained events relevant to this analysis.

a. Energy reconstruction performance

The energy reconstruction is evaluated using simulated electrons with true energies below 10 MeV. Figure 19 compares the reconstructed and true energies after applying the standard fiducial selection. A strong linear correlation demonstrates that the position-dependent calibration provides an unbiased estimate of the deposited energy across the energy range of interest.

The corresponding energy resolution is shown in Fig. 20, where it is defined as the width of the central 68% interval of $(E_{\mathrm{reco}} - E_{\mathrm{true}})$ divided by the true energy. The resolution improves from approximately 12.5% at 1 MeV to 7.5% at 10 MeV, consistent with the expected stochastic scaling of scintillation detectors. This performance provides a reliable energy estimator for the low-energy events considered in the ALP search.

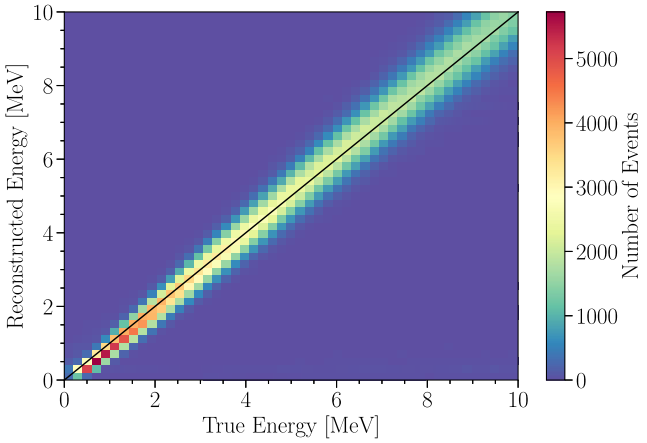


FIG. 19. True versus reconstructed energies for simulated events. The event position is first reconstructed with the GraphNeT model, and the energy is then estimated from the charge recorded in the first 90 ns using the position-dependent calibration factors of Fig. 18.

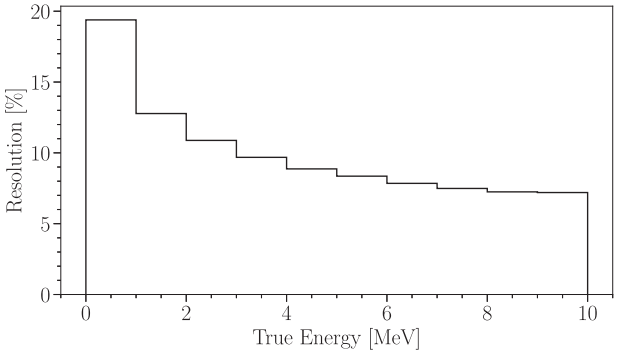


FIG. 20. Energy reconstruction resolution versus true energy. The resolution is evaluated using only particles with true kinetic energies ≤ 10 MeV, corresponding to the energy range relevant for the analysis.

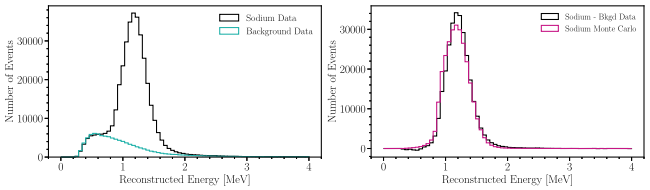


FIG. 21. Energy reconstruction validation with the ^{22}Na calibration source. The left panel compares the reconstructed energy distributions for calibration and background data, while the right panel shows the background-subtracted calibration spectrum overlaid with the Monte Carlo simulation. The reconstructed energy peaks at approximately 1.2 MeV in both data and simulation, with spectral shape differences of less than 10%.

The energy reconstruction is validated using the ^{22}Na calibration dataset. Figure 21 compares the reconstructed energy spectra for source and background data (left), showing a clear excess from ^{22}Na decays. After subtracting the normalized background, the resulting spectrum is compared with the Monte Carlo prediction, as shown in the right panel of Fig. 21.

The reconstructed energy spectra in data and simulation exhibit good agreement in both the energy scale and spectral shape. The peak positions are consistent at approximately 1.2 MeV, while the spectral shapes agree to within ∼10%, consistent with the optical model uncertainties discussed in Ref. [35]. This agreement demonstrates that the position-dependent energy calibration accurately models the detector response and provides a reliable energy estimator for kinematic ranges relevant to this analysis.

[1] R. D. Peccei and H. R. Quinn, Phys. Rev. Lett. **38**, 1440 (1977).
[2] C. A. Baker *et al.*, Phys. Rev. Lett. **97**, 131801 (2006).
[3] C. Abel *et al.*, Phys. Rev. Lett. **124**, 081803 (2020).
[4] J. E. Kim, Phys. Rev. Lett. **43**, 103 (1979).
[5] M. A. Shifman, A. I. Vainshtein, and V. I. Zakharov, Nucl. Phys. **B166**, 493 (1980).
[6] M. Dine, W. Fischler, and M. Srednicki, Phys. Lett. B **104**, 199 (1981).
[7] A. R. Zhitnitsky, Sov. J. Nucl. Phys. **31**, 260 (1980).
[8] G. Lanfranchi, M. Pospelov, and P. Schuster, Annu. Rev. Nucl. Part. Sci. **71**, 279 (2021).
[9] G. Carosi *et al.* (ADMX Collaboration), Phys. Rev. Lett. **135**, 191001 (2025).
[10] K. Altenmüller *et al.* (CAST Collaboration), Phys. Rev. Lett. **133**, 221005 (2024).
[11] E. Armengaud *et al.* (IAXO Collaboration), J. Cosmol. Astropart. Phys. 06 (2019) 047.
[12] X. Bai *et al.* (HAYSTAC Collaboration), Phys. Rev. Lett. **134**, 151006 (2025).
[13] X. Bai *et al.* (HAYSTAC Collaboration), Phys. Rev. D **113**, 012010 (2026).
[14] V. Ankel *et al.*, arXiv:2604.16602.
[15] E. Cortina Gil *et al.* (NA62 Collaboration), Eur. Phys. J. C **85**, 571 (2025).
[16] R. R. Dusaev, D. V. Kirpichnikov, and M. M. Kirsanov, Phys. Rev. D **102**, 055018 (2020).
[17] D. Banerjee *et al.* (NA64 Collaboration), Phys. Rev. Lett. **125**, 081801 (2020).
[18] A. Berlin, S. Gori, P. Schuster, and N. Toro, Phys. Rev. D **98**, 035011 (2018).

[19] N. Blinov, E. Kowalczyk, and M. Wynne, J. High Energy Phys. 02 (2022) 036.
[20] B. J. Park *et al.* (NEON Collaboration), Phys. Rev. Lett. **134**, 201002 (2025).
[21] M. Mirzakhani *et al.*, Phys. Rev. D **112**, 032013 (2025).
[22] G. Alimonti *et al.* (Borexino Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **600**, 568 (2009).
[23] V. Albanese *et al.* (SNO + Collaboration), J. Instrum. **16**, P08059 (2021).
[24] Z. Djurcic *et al.* (JUNO Collaboration), arXiv:1508.07166.
[25] K. Eguchi *et al.* (KamLAND Collaboration), Phys. Rev. Lett. **90**, 021802 (2003).
[26] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. Lett. **135**, 171804 (2025).
[27] M. Askins *et al.* (Theia Collaboration), Eur. Phys. J. C **80**, 416 (2020).
[28] T. Anderson *et al.*, J. Instrum. **18**, P02009 (2023).
[29] M. Ascencio-Sosa *et al.* (ANNIE Collaboration), J. Instrum. **19**, P05070 (2024).
[30] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. D **106**, 012001 (2022).
[31] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. Lett. **129**, 021801 (2022).
[32] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. D **109**, 095017 (2024); **111**, 035030(A) (2025).
[33] B. Dutta, B.-S. Hu, W.-C. Huang, and R. G. Van de Water, Phys. Rev. Lett. **135**, 011801 (2025).
[34] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. D **107**, 095036 (2023).
[35] A. A. Aguilar-Arevalo *et al.* (CCM Collaboration), Phys. Rev. D **112**, 072010 (2025).
[36] D. A. Newmark, First Demonstration of a Hybrid Cherenkov and Scintillation Detector in a Proof-of-Principle Axion Search at a Beam Dump, Ph.D. thesis, MIT, 2026.
[37] P. D. Group and Workman, Prog. Theor. Exp. Phys. **2022**, 083C01 (2022).
[38] A. Thompson and W.-C. Huang, athompson-git/alplib: alpha (2023).
[39] L. Zavorka, M. J. Mocko, and P. E. Koehler, Nucl. Instrum. Methods Phys. Res., Sect. A **901**, 189 (2018).
[40] S. Agostinelli *et al.* (Geant4 Collaboration), Nucl. Instrum. Methods Phys. Res., Sect. A **506**, 250 (2003).
[41] A. Schneider, N. W. Kamp, and A. Y. Wen, Comput. Phys. Commun. **316**, 109799 (2025).
[42] P. Adhikari *et al.* (DEAP Collaboration), Eur. Phys. J. C **81**, 823 (2021).
[43] G. Fiorillo, Nucl. Phys. B, Proc. Suppl. **150**, 372 (2006).
[44] M. Agostini *et al.*, Eur. Phys. J. C **75**, 506 (2015).
[45] C. E. Aalseth *et al.* (DarkSide-20k Collaboration), Eur. Phys. J. Plus **133**, 131 (2018).
[46] C. A. Argüelles, A. Schneider, and T. Yuan, J. High Energy Phys. 06 (2019) 030.
[47] S. S. Wilks, Ann. Math. Stat. **9**, 60 (1938).
[48] J. Jaeckel and M. Spannowsky, Phys. Lett. B **753**, 482 (2016).
[49] J. Blümlein *et al.*, Int. J. Mod. Phys. A **07**, 3835 (1992).
[50] J. Jaeckel, E. Massó, J. Redondo, A. Ringwald, and F. Takahashi, Phys. Rev. D **75**, 013004 (2007).
[51] J. Khoury and A. Weltman, Phys. Rev. Lett. **93**, 171104 (2004).
[52] E. Massó and J. Redondo, Phys. Rev. Lett. **97**, 151802 (2006).
[53] A. Dupays, E. Massó, J. Redondo, and C. Rizzo, Phys. Rev. Lett. **98**, 131802 (2007).
[54] R. N. Mohapatra and S. Nasri, Phys. Rev. Lett. **98**, 050402 (2007).
[55] P. Brax, C. van de Bruck, and A.-C. Davis, Phys. Rev. Lett. **99**, 121103 (2007).
[56] W. DeRocco, P. W. Graham, and S. Rajendran, Phys. Rev. D **102**, 075015 (2020).
[57] E. Massó and J. Redondo, J. Cosmol. Astropart. Phys. 09 (2005) 015.
[58] P. Acosta, J. Bendavid, M. D'Alfonso, J. Eysermans, C. Freer, M. Goncharov *et al.*, EPJ Res. Infrastruct. **2**, 1 (2026).
[59] A. Søgaard *et al.*, J. Open Source Softwaare **8**, 4971 (2023).
[60] H. Bukhari, D. Chakraborty, P. Eller, T. Ito, M. V. Shugaev, and R. Ørsøe, Eur. Phys. J. C **84**, 646 (2024).
[61] D. P. Kingma and J. Ba, arXiv:1412.6980.