

Identifying neutron sources using recoil and time-of-flight spectroscopy

David Breitenmoser¹,* Ricardo Lopez¹, Shaun D. Clarke, and Sara A. Pozzi¹

Department of Nuclear Engineering and Radiological Sciences, *University of Michigan*, 2355 Bonisteel Blvd., Ann Arbor, Michigan 48109-2104, USA



(Received 12 December 2025; revised 16 March 2026; accepted 30 April 2026; published 3 June 2026)

Identification of neutron sources is central to nuclear physics and its applications, from planetary science to nuclear security, yet direct source discrimination from measured neutron spectra remains fundamentally elusive. Here, we introduce a Bayesian protocol that directly infers source ensembles from measured neutron spectra by combining full-spectrum template matching with probabilistic evidence evaluation. Applying this protocol to recoil and time-of-flight spectroscopy, we recover single- and two-source configurations with strong statistical significance (beyond 4σ) at event counts as low as $\sim 10^3$. These results demonstrate that neutron spectral signatures can be leveraged for robust source identification, opening a new observational window for both fundamental research and operationally driven applications.

DOI: [10.1103/v6j6-f4rx](https://doi.org/10.1103/v6j6-f4rx)

I. INTRODUCTION

Neutrons are powerful probes of matter due to their charge neutrality and strong nuclear interactions [1,2]. The energy spectra of neutron fields encode both their production mechanisms and the modulation processes they undergo, making neutron spectroscopy indispensable for applications ranging from determining planetary surface compositions [3–6] to monitoring environmental neutron fields for soil moisture and biomass estimation [7–9]. In safety and security contexts, the ability to infer neutron-source characteristics directly from spectral measurements is equally critical, supporting the discrimination of special nuclear material [10,11] and enabling robust source attribution and risk-informed protective actions in nuclear forensics and emergency response [12,13].

Despite its central role across scientific and applied domains, inferring neutron sources and the modulation processes they undergo directly from measured energy spectra—referred to here as identification—remains fundamentally elusive. Most neutron sources exhibit continuous energy emissions with strong spectral similarity [13,14]. These spectral features, compounded by detector response broadening and the modulating effects of intervening materials, result in ill-conditioned inverse problems that make it difficult to determine the underlying source characteristics

and modulation histories. As a result, current identification methodologies are largely restricted to indirect or proxy signatures, such as secondary alpha and gamma-ray emissions [12,15–18] or correlation metrics derived from integral spectral components [4–6,14]. While secondary emissions can provide additional context, they may be absent or strongly attenuated depending on source encapsulation and environmental conditions. Qualitative spectral metrics, in turn, offer only limited accuracy and precision, are typically restricted to single-source scenarios, and lack quantitative measures of statistical significance.

In this study, we introduce a Bayesian methodology that directly addresses these challenges by enabling quantitative inference of source ensembles using neutron spectroscopy. By combining full-spectrum template matching with probabilistic inference, our approach overcomes the limitations of current methods, providing direct, quantitative discrimination between single- and multisource models, even in the presence of strong spectral similarity and low count statistics. We demonstrate the identification fidelity in a series of dedicated laboratory experiments using varying sets of spontaneous fission and (α ,n) neutron sources, evaluated across two prominent measurement modalities: recoil and time-of-flight (TOF) spectroscopy.

II. EVIDENCE-BASED BAYESIAN SOURCE IDENTIFICATION

Motivated by recent successes in addressing similarly ill-conditioned model inference problems in gravitational-wave astronomy [19–21] and exoplanet research [22,23], we develop here a probabilistic framework for neutron source identification formulated as a Bayesian model comparison problem [24,25]. For a neutron source model $\mathcal{M}$

*Contact author: dbreiten@umich.edu

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license. Further distribution of this work must maintain attribution to the author(s) and the published article's title, journal citation, and DOI.

conditioned on a set of observed neutron spectra $\mathcal{D} = \{\mathbf{y}_i \in \mathbb{N}^N\}_{i=1}^{N_D}$ with measured counts $\mathbf{y} \in \mathbb{N}^N$, $N \in \mathbb{N}$ channels, and $N_D \in \mathbb{N}$ realizations, we quantify the support for $\mathcal{M}$ through the model evidence, defined as the marginal likelihood $p(\mathcal{D} | \mathcal{M})$ given by

$$\mathcal{Z} = \int_{\Theta} \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M}) p(\boldsymbol{\theta} | \mathcal{M}) d\boldsymbol{\theta}, \quad (1)$$

where $\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M})$ and $p(\boldsymbol{\theta} | \mathcal{M})$ denote the likelihood and parameter prior defined on an M -dimensional parameter space $\Theta = \{\boldsymbol{\theta} \in \mathbb{R}^M\}$ with $\mathbb{N}$ and $\mathbb{R}$ representing the sets of natural and real numbers, respectively.

The source model $\mathcal{M}$ represents the forward mapping from any admissible parameter vector $\boldsymbol{\theta}$ to the expected spectral response of a given spectrometric system $\mathcal{M}(\boldsymbol{\theta}) : \Theta \subseteq \mathbb{R}^M \mapsto \mathbb{N}^N$. This mapping is parameterized as a linear superposition of $S \in \mathbb{N}$ spectral templates $\mathcal{M}(\boldsymbol{\theta}) = \sum_{s=1}^S \xi_s \boldsymbol{\psi}_s$, where each template $\boldsymbol{\psi}_s \in \mathbb{R}_+^N$ represents the normalized full-spectrum response of the spectrometric system to the s th neutron source, and $\xi_s \in \mathbb{R}_+$ denotes the corresponding non-negative neutron emission rate. The template set $\mathcal{T} = \{\boldsymbol{\psi}_s\}_{s=1}^S$ thus defines the discrete ensemble of candidate neutron sources considered within a given model $\mathcal{M}$.

To determine which model is most consistent with the data $\mathcal{D}$, we perform Bayesian model comparison across a finite collection of competing models $\{\mathcal{M}_k\}_{k=1}^K$ with associated model priors $p(\mathcal{M}_k)$. The relative support for one model over another is quantified by the posterior odds ratio

$$\mathcal{O}_{ij} = \frac{p(\mathcal{M}_i | \mathcal{D})}{p(\mathcal{M}_j | \mathcal{D})} = \frac{\mathcal{Z}_i}{\mathcal{Z}_j} \frac{p(\mathcal{M}_i)}{p(\mathcal{M}_j)}, \quad (2)$$

which incorporates both the Bayesian evidence $\mathcal{Z}_k$ and the model prior $p(\mathcal{M}_k)$ for model $\mathcal{M}_k$ [26]. In the special case of noncommittal model priors $p(\mathcal{M}_k) = 1/N \forall k$, the posterior odds ratio reduces to the ratio of evidence values, $\mathcal{Z}_i/\mathcal{Z}_j$, which is commonly referred to as the Bayes factor [24,25].

While Eq. (1) is analytically intractable for general likelihood and prior distributions [25], advances in Bayesian computation over the past decades have made accurate numerical approximations feasible. In this work, we employ nested sampling to evaluate Eq. (1) for a predefined set of source models $\{\mathcal{M}_k\}_{k=1}^K$, likelihood functions $\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M})$, and corresponding parameter priors $p(\boldsymbol{\theta} | \mathcal{M})$ [27]. Here, we summarize the likelihood and prior definitions that constitute the Bayesian model, with detailed computational procedures for sampling, convergence diagnostics, and uncertainty quantification provided in Appendix A.

Following analogous treatments in particle physics and cosmology [28–33], we formulate the likelihood $\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M})$ within a probabilistic negative binomial

model $\mathbf{Y} | \boldsymbol{\theta} \sim \mathcal{NB}(\mathbf{y} | \boldsymbol{\theta})$, where $\mathbf{Y}$ denotes the random vector of observed counts. The associated likelihood was formulated using a parameterization established in prior studies [34–36] as

$$\begin{aligned} \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M}) &= \prod_{k=1}^{N_D} \prod_{j=1}^N \frac{\Gamma(y_{k,j} + \frac{1}{\alpha_{\text{NB}}})}{\Gamma(\frac{1}{\alpha_{\text{NB}}}) \Gamma(y_{k,j} + 1)} \\ &\times \left(\frac{1}{1 + \alpha_{\text{NB}} \mathcal{M}(\boldsymbol{\theta})} \right)^{\frac{1}{\alpha_{\text{NB}}}} \\ &\times \left(\frac{1}{1 + [\alpha_{\text{NB}} \mathcal{M}(\boldsymbol{\theta})]^{-1}} \right)^{y_{k,j}} \end{aligned} \quad (3)$$

with $\Gamma(\cdot)$ denoting the gamma function and $y_{k,j} \in \mathbb{N}$ representing the counts observed in channel j of the spectrum k . The dispersion parameter $\alpha_{\text{NB}} \in \mathbb{R}_+$ in Eq. (3) controls the overdispersion beyond the intrinsic Poisson statistics as $\text{Var}(\mathbf{Y}) = \langle \mathbf{Y} \rangle + \alpha_{\text{NB}} \langle \mathbf{Y} \rangle^2$ with $\alpha_{\text{NB}} \in \mathbb{R}_+$. To prevent biased inference, we incorporate the dispersion parameter α_{NB} directly into the inference protocol, estimating it concurrently with the neutron emission rates $\boldsymbol{\xi} \in \mathbb{R}_+^S$. As a result, the inverse problem associated with a given model $\mathcal{M}$ becomes $M = S + 1$ -dimensional, with the full parameter vector defined as $\boldsymbol{\theta} := (\boldsymbol{\xi}, \alpha_{\text{NB}})$.

The likelihood specification in Eq. (3) forms the first component of the Bayesian model definition. Completing the model requires assigning suitable priors to the model parameters $\boldsymbol{\theta}$. To avoid overly restrictive assumptions on these parameter priors, we adopted weakly informative, statistically independent marginal priors $p(\boldsymbol{\theta} | \mathcal{M}) := \prod_{i=1}^M p(\theta_i | \mathcal{M})$ for all model parameters θ , with the marginal priors $p(\boldsymbol{\theta} | \mathcal{M})$ defined based on the principle of maximum entropy [37]. Specifically, we chose a truncated normal distribution for the neutron emission rate $\xi \sim \mathcal{N}(\mu = 10^8 \text{ s}^{-1}, \sigma = 10^8 \text{ s}^{-1}) \in [0, \infty)$, invariant with respect to the source, and an exponential distribution for the dispersion parameter $\alpha_{\text{NB}} \sim \mathcal{E}(\lambda = 1) \in [0, \infty)$.

III. EXPERIMENTS

We demonstrate the fidelity of the proposed Bayesian identification methodology by retrieving spontaneous fission and (α, n) source combinations from observed neutron fields probed by neutron spectroscopy under laboratory conditions.

The neutron fields in all experiments were measured using an array of 12 organic-glass scintillator bars, each coupled to a dual-ended silicon-photomultiplier readout [38], enabling simultaneous recoil and TOF spectroscopy with a single detector system. Three independent experiments were performed covering the complete power set of a $^{252}_{98}\text{Cf}$ and a $^{239}_{94}\text{Pu}$ -Be source, hereafter referred to as Cf-252 and PuBe, with respective neutron emission rates




FIG. 1. Measurement setup for experiment $\mathcal{D}_2$ (single PuBe source at azimuth $\varphi = -35^\circ$ relative to the spectrometer; see Table I).

of $1.74(9) \times 10^6 \text{ s}^{-1}$ and $1.70(4) \times 10^6 \text{ s}^{-1}$. The sources were fixed at a distance of 58.4 cm from the center of the spectrometer's active volume, centered on its horizontal mid-plane. Reproducible placement of sources and the spectrometer was ensured using an optical breadboard and a custom source holder arc (see Fig. 1).

A dedicated background run was performed under identical detector and acquisition conditions to quantify the intrinsic neutron background in both recoil and TOF modes. The background contribution was found to be negligible in all cases, with signal-to-background ratios exceeding 10^4 in recoil spectroscopy and 10^3 in TOF spectroscopy (see Supplemental Material [39]).

Templates ψ for the measured sources were inferred from two additional single-source measurements using a data-driven approach based on generalized additive machine-learning models [41]. To test the robustness and generalization capability of the derived template bank, we deliberately varied the azimuthal angle φ and introduced an additional 6.8 mm lead sleeve during these measurements. A summary of all source-detector configurations and corresponding live times is provided in Table I. Further details on the spectrometer, data reduction pipeline, and template generation protocol can be found in Appendices B–D.

IV. IDENTIFICATION FIDELITY

We report the Bayesian evidence results for the complete set of competing models, defined as the power set of single-source templates, $\mathcal{P}(\mathcal{S}) = \{\mathcal{M}_k\}_{k=1}^K$, where $\mathcal{S}$ denotes the set of individual single-source templates. Results are shown for the three identification experiments in Figs. 2 and 3 using recoil and TOF spectroscopy. To




FIG. 2. Bayesian evidence results for three experiments using recoil spectroscopy. (a),(b) Single-source Cf-252 experiment ($\mathcal{D}_1$). (c),(d) Single-source PuBe experiment ($\mathcal{D}_2$). (e),(f) Two-source Cf-252 and PuBe experiment ($\mathcal{D}_3$). Panels (a),(c),(e) show log evidence values $\log \mathcal{Z}_{i=j}$ for each model $\mathcal{M}_{i=j}$ (antidiagonal entries), log Bayes factors $\Delta \log \mathcal{Z}_{ij} = \log(\mathcal{Z}_i/\mathcal{Z}_j)$ (above antidiagonal entries), and log posterior odds ratios $\log \mathcal{O}_{ji} = \log(\mathcal{Z}_j/\mathcal{Z}_i) + \log[p(\mathcal{M}_j)/p(\mathcal{M}_i)]$ (below antidiagonal entries) with $p(\mathcal{M}) \propto 4^{-\dim(\mathcal{M})}$. Uncertainties are indicated using least-significant-figure notation, with lower statistical-significance bounds in square brackets [24]. Panels (b),(d),(f) show the measured energy spectra (coverage factor $k=1$) alongside maximum-a-posteriori predictions and 99% central posterior predictive intervals (shaded area) for the retrieved (true) source set.

illustrate the robustness of our results across retrieval assumptions, we explore the model-selection fidelity using both log Bayes factors $\Delta \log \mathcal{Z}$, representing noncommittal model prior assumptions, as well as log posterior odds ratios $\log \mathcal{O}$, adopting weakly informative model priors $p(\mathcal{M}) \propto 4^{-\dim(\mathcal{M})}$ favoring lower-dimensional models [26,42]. For both spectroscopy modalities, the true source set is recovered with decisive statistical confidence on Jeffreys' scale ($\Delta \log \mathcal{Z} > 8.0$, $\log \mathcal{O} > 8.7$) over all competing models [43], corresponding to a statistical significance beyond 4σ [24].

TABLE I. Key measurement parameters of the performed neutron spectroscopy experiments.

ID	Source	ξ (s ⁻¹) ^a	φ (°) ^b	t_{live} (s) ^c	Casing ^d
$\mathcal{D}_1$	Cf-252	$1.74(9) \times 10^6$	+35	5.63×10^4	✗
$\mathcal{D}_2$	PuBe	$1.70(4) \times 10^6$	-35	1.80×10^4	✗
$\mathcal{D}_3$	Cf-252	$1.74(9) \times 10^6$	+35	2.03×10^4	✗
	PuBe	$1.70(4) \times 10^6$	-35		✗
$\mathcal{D}_4$	Cf-252	$1.74(9) \times 10^6$	0	5.63×10^4	✓
$\mathcal{D}_5$	PuBe	$1.70(4) \times 10^6$	0	1.80×10^4	✓
$\mathcal{D}_b$	background			5.32×10^4	

^a Neutron emission rate.^b Azimuth source position with respect to the spectrometer (see Fig. 1).^c Measurement live time (rounded to three significant figures).^d Flag indicating use of a lead sleeve (6.8 mm thick) around the deployed neutron source(s) (see Fig. 1).

Comparing noncommittal and weakly informative model priors, we observe statistically significant differences in the inferred evidence for the true source configuration, which nevertheless do not alter the model-selection outcome. For example, in the recoil spectroscopy modality of the Cf-252 and PuBe experiment [Fig. 2(e)], the log posterior odds ratio (weakly informative prior) is 106.1(2), while the corresponding log Bayes factor (noncommittal

prior) is 106.8(2) relative to the best competing model (PuBe), illustrating that the model prior choice has a negligible effect on the model selection fidelity, and confirming that the outcome is dictated by the observed data rather than by the specific form of the model priors for the given measurement setup.

Here, all identification results were obtained with a fixed measurement time $\mathcal{O}(10^4)$ s per experiment, with recoil and TOF spectroscopy acquisitions performed in parallel. Because TOF spectroscopy relies on coincidence events, its event rate is $\mathcal{O}(10^2)$ lower than that of recoil spectroscopy, leading to consistently higher log Bayes factors and log posterior odds ratios for the recoil modality. To generalize beyond system-specific event rates, we quantify in the next section the scaling of model-selection fidelity with the total number of detected events for both spectroscopy modalities.

V. EVENT SCALING

To assess the scaling of model-selection fidelity with the total number of detected events, we compute the posterior probability $p(\mathcal{M}_{\text{true}} | \mathcal{D})$ of retrieving the true source model $\mathcal{M}_{\text{true}}$ conditioned on the experimental data $\mathcal{D}$ for varying event counts (see Appendix A). Additionally, to explore the impact of the applied prior assumptions on $p(\mathcal{M}_{\text{true}} | \mathcal{D})$, we relax the model prior used previously to $p(\mathcal{M}) \propto 2^{\lambda \dim(\mathcal{M})}$ with $\lambda \in [-3, +3]$ and $\lambda = 0$ representing a noncommittal model prior. In Fig. 4, the results of these computations for recoil and TOF spectroscopy are presented. As expected, the experiments differ in how quickly they decisively favor the true model, with the single-source cases achieving $p(\mathcal{M}_{\text{true}} | \mathcal{D}) \approx 1$ after $\sim 10^2$ events, whereas the two-source cases require $\sim 10^3$ events. Owing to the natural regularization inherent in the Bayesian evidence, we observe that increasing the model prior exponent λ consistently leads to faster convergence by favoring higher-dimensional models, while decreasing it slows convergence by favoring lower-dimensional models, consistent with the dimensionality weighting in the




FIG. 3. Same as Fig. 2 but using TOF instead of recoil spectroscopy.




FIG. 4. Posterior model probability $p(\mathcal{M}_{\text{true}} | \mathcal{D})$ of retrieving the true source model $\mathcal{M}_{\text{true}}$ as a function of the total number of detected events, for varying prior exponents λ in the model prior $p(\mathcal{M}) \propto 2^{\lambda \dim(\mathcal{M})}$ with $\lambda \in [-3, +3]$. Panels (a),(b) show the single-source Cf-252 experiment ($\mathcal{D}_1$), (c),(d) the single-source PuBe experiment ($\mathcal{D}_2$), and (e),(f) the two-source Cf-252 and PuBe experiment ($\mathcal{D}_3$). Panels (a),(c),(e) display posteriors conditioned on recoil spectroscopy data, while (b),(d),(f) show posteriors conditioned on TOF spectroscopy data. In addition to the number of detected events, the corresponding measurement times for each spectroscopy modality are indicated.

prior. Interestingly, recoil spectroscopy converges consistently faster than TOF spectroscopy, not only in temporal terms, as discussed above, but also with respect to the absolute event count. This difference in scaling motivates an information-theoretic analysis of the two modalities to determine whether one yields an inherently higher information gain than the other on an absolute event scale.

Following an information-theoretic perspective [44], we aim to quantify the information gained $\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D})$ about the true source model $\mathcal{M}_{\text{true}}$ from our experimental data $\mathcal{D}$ for varying event counts relative to our prior beliefs $p(\mathcal{M}_{\text{true}})$ over all candidate models. Building on earlier information theory work in Bayesian parameter inference [44–46], we define this information gain as the relative entropy [47], which measures the information (in bits) gained about the true source model $\mathcal{M}_{\text{true}}$ from the dataset $\mathcal{D}$ (see Appendix A). In the limit where




FIG. 5. Information gain $\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D})$ about the true source model $\mathcal{M}_{\text{true}}$ from the dataset $\mathcal{D}$ as a function of the total number of detected events, for varying prior exponents λ in the model prior $p(\mathcal{M}) \propto 2^{\lambda \dim(\mathcal{M})}$ with $\lambda \in [-3, +3]$. Panels (a),(b) show the single-source Cf-252 experiment ($\mathcal{D}_1$), (c),(d) the single-source PuBe experiment ($\mathcal{D}_2$), and (e),(f) the two-source Cf-252 and PuBe experiment ($\mathcal{D}_3$). Panels (a),(c),(e) display $\mathcal{IG}$ for recoil spectroscopy data, while (b),(d),(f) show $\mathcal{IG}$ for TOF spectroscopy data. In addition to the number of detected events, the corresponding measurement times for each spectroscopy modality are indicated.

the posterior collapses on the true model $p(\mathcal{M}_{\text{true}} | \mathcal{D}) \rightarrow 1$, $\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D})$ approaches $-\log_2 p(\mathcal{M}_{\text{true}})$. Thus, the prior places an absolute upper bound on the achievable information. In Fig. 5, we show $\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D})$ for all three experiments as a function of event count. Across all tested prior weights and source configurations, recoil spectroscopy exhibits a consistently steeper rise in information gain than TOF spectroscopy, demonstrating that—beyond the expected advantage from its higher event rate—recoil spectroscopy also achieves a higher information-acquisition rate on an absolute event scale. Based on both parameter and predictive posteriors across all experiments, we attribute the lower information gain in TOF spectroscopy to its substantially higher overdispersion relative to recoil spectroscopy (factor $\sim 10^1$ difference in the dispersion parameter α_{NB}), resulting in an inherently




FIG. 6. Posterior model probability $p(\mathcal{M}_{\text{true}} | \mathcal{D})$ of retrieving the true source model $\mathcal{M}_{\text{true}}$ (Cf-252 and PuBe) as a function of the total number of detected events, for varying neutron emission ratios $\xi_{\text{Cf-252}}/\xi_{\text{PuBe}}$ with noncommittal model priors $p(\mathcal{M}) = \text{const.}$ (corresponding to $\lambda = 0$). Panel (a) displays posteriors conditioned on recoil spectroscopy data, while (b) shows posteriors conditioned on TOF spectroscopy data. For reference, we indicate, in addition to the number of detected events, measurement times for each spectroscopy modality assuming a neutron emission rate $\xi_{\text{PuBe}} = 1.70(4) \times 10^6 \text{ s}^{-1}$ for PuBe reflective of the one used in our experiments (see Sec. III). Uncertainties are indicated as 3σ shaded areas.

higher statistical dispersion per event. Full posterior results for all experiments are provided in the Supplemental Material [39].

We conclude the event-scaling analysis by examining how the posterior probability $p(\mathcal{M}_{\text{true}} | \mathcal{D})$ of retrieving the true source model $\mathcal{M}_{\text{true}}$ scales with the total number of detected events when the neutron emission ratio in our experiments $\xi_{\text{Cf-252}}/\xi_{\text{PuBe}} \approx 1$ is relaxed (see Sec. III). To this end, we generated two-source datasets by combining resampled Cf-252 and PuBe single-source spectroscopic data with varying emission ratios using a resampling protocol described in Appendix A, in direct analogy to the analysis of Fig. 4. The resulting posterior probabilities are shown in Fig. 6. As expected, greater imbalance between the neutron emission rates requires more events for the posterior to decisively favor the true model. Nevertheless, even under significant imbalances $\xi_{\text{Cf-252}}/\xi_{\text{PuBe}} = \mathcal{O}(10^2)$, the correct source set is recovered with decisive statistical confidence on Jeffreys' scale ($\Delta \log \mathcal{Z} > 13$, exceeding 5σ) after $\mathcal{O}(10^6)$ detected events for both spectroscopy modalities, demonstrating the robustness of the Bayesian identification procedure even under strongly nonuniform source contributions. Consistent with the findings above, recoil spectroscopy reaches decisive posterior probabilities at systematically lower event counts than TOF spectroscopy, indicating that its advantage persists even when the constraint of comparable neutron emission rates is relaxed.

VI. CONCLUSION

In this study, we have demonstrated the first quantitative identification of neutron sources from measured neutron spectra by combining full-spectrum template matching with probabilistic evidence evaluation. Our framework achieves high-confidence discrimination (beyond 4σ) in single- and two-source configurations with comparable neutron emission rates at event counts as low as 10^3 . Furthermore, we find that this identification remains robust to source imbalances: even at neutron emission ratios of $\mathcal{O}(10^2)$, the true source model is retrieved with decisive statistical confidence within $\mathcal{O}(10^6)$ events. Comparing recoil and TOF modalities, we find that recoil spectroscopy delivers consistently higher information gain across all source configurations, both on a temporal and on a per-event basis. This work demonstrates that neutron spectral signatures can be leveraged as a direct and quantitative modality for source identification, opening a new, robust observational window for both basic research and field applications.

While our study focused on simplified source sets under controlled laboratory conditions, the framework naturally extends to more complex and realistic source scenarios. To scale these methods efficiently without the need for prohibitively large measurement campaigns, future work will explore the use of high-fidelity neutron transport simulations [48–50] to generate comprehensive, physics-informed template banks. By incorporating spectral variability arising from different geometries, shielding materials, or environmental conditions, such physics-informed modeling will enable the joint quantitative inference of both the source ensembles and the modulation processes the neutrons undergo. This capability will extend our retrieval methods to a broad variety of applications, from constraining the surface composition of planetary bodies in the solar system [3–6] to supporting discrimination of special nuclear material [10,11] and robust source attribution in nuclear forensics, radiochemistry, and emergency response [12,13].

ACKNOWLEDGMENTS

This work was funded in-part by the Consortium for Monitoring, Technology, and Verification under Department of Energy National Nuclear Security Administration Award No. DE-NA0003920.

DATA AVAILABILITY

The data that support the findings of this article are openly available [51]; embargo periods may apply.

APPENDIX A: BAYESIAN COMPUTATIONS

In this appendix, we describe the numerical evaluation of the Bayesian evidence in Eq. (1) using nested sampling alongside a thorough discussion of the convergence diagnostics and uncertainty quantification. We also provide details on the Bayesian computations for the event-scale analysis in Sec. V, including the evaluation of posterior model probabilities and information-gain measures.

1. Bayesian evidence

Using the specified joint prior distribution and likelihood function in Sec. II, we evaluated the Bayesian evidence $\mathcal{Z}$ in Eq. (1) while simultaneously sampling the posterior distribution $p(\boldsymbol{\theta} | \mathcal{D}, \mathcal{M})$ using the nested sampling code `dynesty` (version 3.0.0) [52]. This Bayesian inference procedure was carried out independently for each candidate source model $\{\mathcal{M}_k\}_{k=1}^K$, yielding a corresponding set of Bayesian evidence values $\{\mathcal{Z}_k\}_{k=1}^K$ from which Bayes factors and posterior odds ratios are readily obtained as detailed in Sec. II. For each nested-sampling run, we used 2^{10} live points and set the evidence tolerance to $\Delta \log \mathcal{Z} = 0.1$. The bounding strategy `single` was adopted, and new samples were drawn using the default uniform-sampling method. All other sampler settings were kept at their default values. Convergence and numerical precision were evaluated using established nested-sampling diagnostics [27,52,53]. Across all runs, we obtained an effective sample size $\text{ESS} > 4 \times 10^3$ and a sampling efficiency exceeding 25%. Following the uncertainty quantification framework introduced in Refs. [53–55], the statistical uncertainties on the computed Bayesian evidence values were estimated directly from the `dynesty` outputs and are reported in Sec. IV. In addition to the Bayesian evidence values, Bayes factors, and posterior odds ratios reported in Sec. IV, we provide the full posterior results for all nested sampling runs in the Supplemental Material [39].

It is important to note that the gamma function terms in Eq. (3) increase rapidly for moderately large arguments, requiring careful numerical evaluation of the likelihood function. To prevent overflow, we conducted all Bayesian computations using the logarithm of the likelihood function:

$$\begin{aligned} \log \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M}) &= \sum_{k=1}^{N_{\mathcal{D}}} \sum_{j=1}^N \log \Gamma\left(y_{k,j} + \frac{1}{\alpha_{\text{NB}}}\right) \\ &\quad - \log \Gamma\left(\frac{1}{\alpha_{\text{NB}}}\right) - \log \Gamma(y_{k,j} + 1) \\ &\quad - \frac{1}{\alpha_{\text{NB}}} \log(1 + \alpha_{\text{NB}} \mathcal{M}(\boldsymbol{\theta})) \\ &\quad - y_{k,j} \log\left(1 + \frac{1}{\alpha_{\text{NB}} \mathcal{M}(\boldsymbol{\theta})}\right). \end{aligned} \quad (\text{A1})$$

The log-gamma terms in Eq. (A1) were computed using a numerically robust algorithm [56,57], which maintains stability and precision across the full dynamic range of our input arguments.

2. Posterior model probability

In Sec. V, we evaluated how the posterior model probability $p(\mathcal{M}_{\text{true}} | \mathcal{D})$ and the information gain $\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D})$ relative to the true source model $\mathcal{M}_{\text{true}}$ scale with the total number of detected events in the experiments performed. In the following two sections, we detail the computational procedure used to obtain these quantities.

We started our event-scale analysis by generating a collection of resampled datasets $\{\mathcal{D}^{(b)} = \{\mathbf{y}^{(b)} \in \mathbb{N}^N\}\}_{b=1}^B$, where each spectrum $\mathbf{y}^{(b)}$ contained exactly $N_{\text{event}}^{(b)}$ total counts drawn from a predefined reference dataset $\mathcal{D}$ corresponding to one of the recorded spectra $\mathbf{y} \in \mathbb{N}^N$ presented in Sec. IV. Defining the spectral probability vector $\mathbf{p} = (p_1, \dots, p_N)$ with components $p_i = y_i / \sum_{j=1}^N y_j$, each spectrum was generated via multinomial sampling $\mathbf{y}^{(b)} \sim \text{Multinomial}(N_{\text{event}}^{(b)}, \mathbf{p}) \in \mathbb{N}^N$ with probabilities $\mathbf{p}$ proportional to the reference spectrum $\mathbf{y}$. This procedure yielded independent finite-statistics realizations that preserve the spectral shape of $\mathbf{y}$. For each number of events $N_{\text{event}}^{(b)}$, we generated $N_{\text{rep}} = 10^3$ realizations, with $N_{\text{event}}^{(b)}$ defined on a logarithmically spaced grid with 4×10^2 instances between $N_{\text{event}}^{(1)} = 1$ and $N_{\text{event}}^{(B)} = N_{\text{event}}$, with N_{event} being the number of detected events in the reference spectrum $\mathbf{y}$.

We continued by computing the log Bayesian evidence $\log \mathcal{Z}^{(b)}$ for each resampled dataset $\mathcal{D}^{(b)}$ as

$$\log \mathcal{Z}^{(b)} = \log \int_{\Theta} \mathcal{L}(\boldsymbol{\theta}; \mathcal{D}^{(b)}, \mathcal{M}) p(\boldsymbol{\theta} | \mathcal{M}) d\boldsymbol{\theta} \quad (\text{A2a})$$

$$= \log \mathcal{Z} - \log \int_{\Theta} \underbrace{\frac{\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}^{(b)}, \mathcal{M})}{\mathcal{L}(\boldsymbol{\theta}; \mathcal{D}, \mathcal{M})}}_{\text{likelihood ratio}} \underbrace{p(\boldsymbol{\theta} | \mathcal{D}, \mathcal{M})}_{\text{posterior}} d\boldsymbol{\theta} \quad (\text{A2b})$$

$$\begin{aligned} &\approx \log \mathcal{Z} - \log N_{\theta} + \log \sum_{i=1}^{N_{\theta}} \exp [\\ &\quad \log \mathcal{L}(\boldsymbol{\theta}_i; \mathcal{D}^{(b)}, \mathcal{M}) - \log \mathcal{L}(\boldsymbol{\theta}_i; \mathcal{D}, \mathcal{M})] \end{aligned} \quad (\text{A2c})$$

with $\log \mathcal{Z}$ denoting the log Bayesian evidence of the reference dataset $\mathcal{D}$, while $\log \mathcal{L}(\boldsymbol{\theta}_i; \mathcal{D}, \mathcal{M})$ and $\log \mathcal{L}(\boldsymbol{\theta}_i; \mathcal{D}^{(b)}, \mathcal{M})$ are the log likelihoods for $\mathcal{D}$ and $\mathcal{D}^{(b)}$ evaluated at each posterior sample $\{\boldsymbol{\theta}_i\}_{i=1}^{N_{\theta}}$ obtained from the nested sampling runs (see Appendix A 1). The importance sampling procedure in Eq. (A2c) allowed us to efficiently estimate the log Bayesian evidence for each of the 400×1000 resampled spectra, from which we then

computed the mean and standard error using Monte Carlo error propagation across the $N_{\text{rep}} = 10^3$ realizations.

Based on the computed set of log Bayesian evidence values $\log \mathcal{Z}^{(b)}$ for each resampled dataset $\mathcal{D}^{(b)}$, we computed the posterior model probability, that is, the posterior probability $p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)})$ of retrieving the true source model $\mathcal{M}_{\text{true}}$ conditioned on $\mathcal{D}^{(b)}$ and a set of competing models $\{\mathcal{M}_k\}_{k=1}^K$ with associated model priors $p(\mathcal{M}_k)$:

$$p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)}) = \frac{p(\mathcal{D}^{(b)} | \mathcal{M}_{\text{true}}) p(\mathcal{M}_{\text{true}})}{\sum_{k=1}^K p(\mathcal{D}^{(b)} | \mathcal{M}_k) p(\mathcal{M}_k)} \quad (\text{A3})$$

with $p(\mathcal{D}^{(b)} | \mathcal{M}_k) = \mathcal{Z}_k^{(b)}$ denoting the Bayesian evidence of $\mathcal{D}^{(b)}$ conditioned on model $\mathcal{M}_k$. For numerical stability, Eq. (A3) was evaluated in logarithmic form using the Log-SumExp (LSE) transformation, which allows for robust computation of the logarithm of a sum of exponentials and prevents over- or underflow:

$$\begin{aligned} \log p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)}) &= \log \mathcal{Z}_{\text{true}}^{(b)} + \log p(\mathcal{M}_{\text{true}}) \\ &\quad - \log \sum_{k=1}^K \exp \left[\log \mathcal{Z}_k^{(b)} + \log p(\mathcal{M}_k) \right] \end{aligned} \quad (\text{A4})$$

where $\mathcal{Z}_{\text{true}}^{(b)} := p(\mathcal{D}^{(b)} | \mathcal{M}_{\text{true}})$ denotes the Bayesian evidence of $\mathcal{D}^{(b)}$ conditioned on the true model $\mathcal{M}_{\text{true}}$. As with the Bayesian evidence itself, the statistical uncertainty in the posterior model probability $p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)})$ was quantified via Monte Carlo error propagation over the $N_{\text{rep}} = 10^3$ available realizations.

The computational protocol described above was further generalized to evaluate the source configurations presented in Fig. 6. For these evaluations, resampled datasets $\mathcal{D}^{(b)}$ were constructed by combining single-source spectral components with prescribed neutron emission ratios $\xi_{\text{Cf-252}}/\xi_{\text{PuBe}} \in \{1, 10^1, 10^2\}$. To ensure robust resampling across the extended dynamic range of event counts, the multinomial sampling probabilities $\mathbf{p}$ were defined proportional to the high-fidelity spectral templates Ψ (see Sec. II and Appendix D). For each configuration, we utilized $N_{\text{rep}} = 10^3$ realizations per grid point, with total event counts $N_{\text{event}}^{(b)}$ distributed across a logarithmically spaced grid of 10^2 instances from $N_{\text{event}}^{(1)} = 1$ to a maximum of $N_{\text{event}}^{(B)} = (1 + \xi_{\text{Cf-252}}/\xi_{\text{PuBe}})10^8$. Bayesian evidence values for these realizations were computed via Eq. (A2c), where the reference log-evidence $\log \mathcal{Z}$ and log-likelihood $\log \mathcal{L}(\theta; \mathcal{D}, \mathcal{M})$ were defined relative to the high-statistics dataset $\mathcal{D} := \mathcal{D}^{(B)}$ and retrieved using the nested-sampling procedure detailed in Sec. A 1. The statistical uncertainty in the posterior model probability $p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)})$ was then again quantified via Monte Carlo error propagation over the available realizations.

3. Information gain

Building on earlier information theory work in Bayesian parameter inference [44–46], we continued our event-scale analysis by assessing the information gained about the true source model $\mathcal{M}_{\text{true}}$ from the dataset $\mathcal{D}^{(b)}$. For each resampled dataset $\mathcal{D}^{(b)}$, we quantified the information gain as the Kullback-Leibler divergence (relative entropy) between the posterior and prior probabilities for the true model [47]:

$$\mathcal{IG}(\mathcal{M}_{\text{true}}; \mathcal{D}^{(b)}) = p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)}) \log_2 \frac{p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)})}{p(\mathcal{M}_{\text{true}})}, \quad (\text{A5})$$

where $p(\mathcal{M}_{\text{true}} | \mathcal{D}^{(b)})$ is the posterior model probability obtained through Eq. (A4). This formulation quantifies the information gained about the true model from the observed data relative to the prior expectation in units of bits. The statistical uncertainty in $\mathcal{IG}$ was again estimated by propagating the Monte Carlo variability across the $N_{\text{rep}} = 10^3$ realizations of the resampled datasets. Averaging over these realizations yielded the expected information gain and its standard error, providing a robust measure of how strongly each dataset supports the true model under the competing hypothesis set $\{\mathcal{M}_k\}_{k=1}^K$.

APPENDIX B: NEUTRON SPECTROMETER

The neutron spectrometer employed in this work is an organic-glass scintillator (OGS) based system that has been characterized extensively in prior studies [14,38,58]. OGS material offers a combination of high light output, fast scintillation response, and excellent pulse shape discrimination (PSD), making it particularly well suited for neutron spectroscopy across both recoil and TOF modalities [59–61]. The OGS composition used in this work consists of a 90 : 10 molecular blend of bis(9,9-dimethyl-9H-fluoren-2-yl)diphenylsilane, $\text{C}_{42}\text{H}_{36}\text{Si}$, and phenyltris(9,9-dimethyl-9H-fluorene-2-yl)silane, $\text{C}_{51}\text{H}_{44}\text{Si}$, combined with a 0.2 wt% 1,4-bis(2-methylstyryl)benzene, $\text{C}_{24}\text{H}_{22}$, wavelength shifter (Bis-MSB). Detailed information about the chemistry and synthesis of these compounds can be found in Ref. [62].

The spectrometer consists of 12 OGS bars, each with dimensions of $6 \times 6 \times 50 \text{ mm}^3$ (see Fig. 1). All bars are wrapped in polytetrafluoroethylene (Teflon) tape to enhance diffuse optical reflection and improve light-collection uniformity. The OGS bars are coupled between two 64-pixel ArrayJ-60035-64P-PCB Onsemi silicon photomultiplier (SiPM) boards operated at 29.6 V. This dual-ended readout configuration allows for event reconstruction on the vertical axis of the OGS bars [38,63]. Prior work has demonstrated a time resolution of 270 ps at 200–341 keV between two OGS bars [64]. The spectrometer also features eight encapsulated cylindrical CeBr_3

crystals (6 mm diameter and height) for gamma-ray spectroscopy and imaging, which are not relevant for this work. All scintillator/SiPM interfaces are coupled using 0.5-mm-thick silicone rubber EJ-560 pads to ensure stable optical transmission. The system uses custom-made printed circuit boards (PCBs) to apply the voltage bias and readout the signals from the SiPM arrays [65]. The signals from all channels are digitized using synchronized CAEN v1730 waveform digitizers (14-bit resolution, 500 MHz sampling rate, 2 V dynamic range). The CAEN v1730 modules were synchronized using a shared reference clock to maintain consistent timing across boards, enabling accurate reconstruction of interbar time differences for TOF analysis. The full detector assembly is enclosed in an aluminum dark box to suppress ambient light. A schematic photograph of the instrument and its mechanical layout is shown in Fig. 1.

APPENDIX C: DATA REDUCTION

In this work, we adopt a benchmarked data reduction pipeline that has already been discussed in detail in our previous work [38,66]. In this subsection, we limit the discussion to key steps used for recoil and TOF spectroscopy. Further information on waveform processing, filtering, and uncertainty quantification is provided in Refs. [38,66].

Data are collected in full-waveform acquisition mode with a record length of 800 ns. Energy calibration was performed using ^{137}Cs measurements at the start and end of each experiment, following a calibration protocol described in Ref. [67]. For each detected event, the signals recorded at each end of the organic-glass scintillator bar are summed together if triggered within a 20 ns coincidence window. The start time of each detected event is determined by the average of the start times from the paired signals. Neutron and gamma-ray signals were separated using charge-integration PSD with an autoslicing algorithm [68,69]. This PSD step is applied to all events and is common to both recoil and TOF spectroscopy.

1. Recoil spectroscopy

In recoil spectroscopy mode, we process all PSD-classified neutron events from individual OGS bars. For each event, the scintillation pulse produced by the neutron-induced proton recoil in OGS is integrated to obtain the pulse integral, which quantifies the light output. The calibrated light output is then expressed in electron-equivalent deposited energy to generate the recoil energy spectra.

2. TOF spectroscopy

TOF spectroscopy extends the analysis to multibar coincidence events to reconstruct the incident neutron energy. The procedure begins with the same PSD-classified neutron population used for recoil spectroscopy. From this set, candidate TOF events are selected by requiring a two-bar

coincidence within a timing window of 0.137–9 ns, motivated by the geometrical constraints of the array and the selected neutron energy band. This filter isolates neutrons that undergo sequential inelastic scattering in two different bars while suppressing accidental coincidences. For each accepted coincidence event, the incident neutron energy E_0 is then reconstructed as

$$E_0 = E_{\text{dep},1} + E_{\text{TOF}} \quad (\text{C1})$$

where $E_{\text{dep},1}$ denotes the energy deposited in the first scatter. This quantity is derived from the pulse integral light output and corrected for the nonproportional scintillation response of OGS using the Birks model [58,70]. The second term, E_{TOF} , represents the neutron energy inferred from the measured interbar flight time and the reconstructed interaction positions along each bar [38,66]. To mitigate systematic biases from gamma-ray leakage, pileup, and geometrically inconsistent scatter sequences, we impose an additional kinematic constraint

$$E_{\text{dep},2} \leq E_{\text{TOF}} \quad (\text{C2})$$

where $E_{\text{dep},2}$ is the energy deposited in the second scatter. This requirement enforces that the second interaction cannot deposit more energy than is kinematically available from the neutron after the first interaction. Finally, we exclude events for which the relative uncertainty in the reconstructed neutron energy exceeds 50%, ensuring a high-quality dataset dominated by well-reconstructed trajectories. The reconstructed set of incident neutron energies $\{E_0\}$ is then used to generate the TOF neutron spectra presented in Sec. IV.

APPENDIX D: TEMPLATE GENERATION

As discussed in Sec. II, we define the source model $\mathcal{M}$ in this work as the discrete ensemble of $S \in \mathbb{N}$ candidate neutron sources together with their associated template set $\mathcal{T} = \{\psi_s\}_{s=1}^S$. Each template $\psi_s \in \mathbb{R}_+^N$ represents the normalized full-spectrum response of the spectrometric system with $N \in \mathbb{N}$ channels to the s th neutron source. Following a full-spectrum template matching approach, we then parameterize the forward mapping from any admissible parameter vector $\theta \in \mathbb{R}^M$ to the expected spectral response of a given spectrometric system $\mathcal{M}(\theta) : \Theta \subseteq \mathbb{R}^M \mapsto \mathbb{R}^N$ as the linear superposition of $S \in \mathbb{N}$ spectral templates $\mathcal{M}(\xi) = \sum_{s=1}^S \xi_s \psi_s$, where $\xi_s \in \mathbb{R}_+$ denotes the corresponding neutron emission rate.

To generate the template bank required to evaluate the full power set of a Cf-252 and PuBe source considered in Sec. IV, we performed two additional single-source Cf-252 ($\mathcal{D}_4$) and PuBe ($\mathcal{D}_5$) calibration experiments, as discussed in Sec. III. Adopting a data-driven strategy to minimize model bias, we generated the templates using a generalized




FIG. 7. Generalized-additive-model-based spectral templates for Cf-252 and PuBe together with the corresponding calibration data ($\mathcal{D}_4$ and $\mathcal{D}_5$; see Table I) for two spectroscopy modalities: (a) recoil spectroscopy; (b) TOF spectroscopy. Measurement uncertainties are shown with coverage factor $k = 1$, while template uncertainties are indicated as 99% prediction intervals. For better interpretability, the templates are scaled by the neutron emission rate of each source to represent absolute neutron counts.

additive model (GAM) implemented via the `pyGAM` package [41]. For each source, we inferred a semiparametric GAM from the calibration data using penalized iteratively reweighted least squares (10^{-5} tolerance). We employed a variable number of F penalized B-spline basis functions ($F \in [30, 70]$ for recoil and $F \in [2, 12]$ for TOF spectroscopy), a Poisson log-link, and a regularization strength $\lambda_{\text{GAM}} \in [10^{-1}, 10^3]$. Both the number of basis functions F and the regularization parameter λ_{GAM} were selected through hyperparameter tuning. The learned spectral templates, together with the underlying calibration data, are shown in Fig. 7 for recoil and TOF spectroscopy.

As discussed in Sec. III, the source position and source casing were deliberately varied between the template generation and identification experiments to demonstrate the robustness and generalization capability of the derived templates. To avoid any significant bias in the retrieval of the source set between the templates and the experimental data, we restricted the energy bands used for identification to 0.1–3.5 MeVee for recoil spectroscopy (bin width $\Delta E = 0.05$ MeVee, $N = 68$ channels) and 2–10 MeV for TOF spectroscopy (bin width $\Delta E = 0.1$ MeV, $N = 80$ channels). A sensitivity study confirmed that the identification results are stable under 10 keV variations of these band limits (see Supplemental Material [39]).

- [2] A. Pietropaolo, M. Angelone, R. Bedogni, N. Colonna, A. J. Hurd, A. Khaplanov, F. Murtas, M. Pillon, F. Piscitelli, E. M. Schooneveld, and K. Zeitelhack, Neutron detection techniques from μeV to GeV, *Phys. Rep.* **875**, 1 (2020).
- [3] I. Mitrofanov, D. Anfimov, A. Kozyrev, M. Litvak, A. Sanin, V. Tret'yakov, A. Krylov, V. Shvetsov, W. Boynton, C. Shinohara, D. Hamara, and R. S. Saunders, Maps of sub-surface hydrogen from the High Energy Neutron Detector, Mars Odyssey, *Science* **297**, 78 (2002).
- [4] W. C. Feldman, S. Maurice, A. B. Binder, B. L. Barraclough, R. C. Elphic, and D. J. Lawrence, Fluxes of fast and epithermal neutrons from lunar prospector: Evidence for water ice at the lunar poles, *Science* **281**, 1496 (1998).
- [5] T. H. Prettyman, D. W. Mittlefehldt, N. Yamashita, D. J. Lawrence, A. W. Beck, W. C. Feldman, T. J. McCoy, H. Y. McSween, M. J. Toplis, T. N. Titus, P. Tricarico, R. C. Reedy, J. S. Hendricks, O. Forni, L. Le Corre, J. Y. Li, H. Mizzon, V. Reddy, C. A. Raymond, and C. T. Russell, Elemental mapping by Dawn reveals exogenic H in Vesta's regolith, *Science* **338**, 242 (2012).
- [6] D. J. Lawrence, W. C. Feldman, J. O. Goldsten, S. Maurice, P. N. Peplowski, B. J. Anderson, D. Bazell, R. L. McNutt, L. R. Nittler, T. H. Prettyman, D. J. Rodgers, S. C. Solomon, and S. Z. Weider, Evidence for water ice near Mercury's north pole from MESSENGER neutron spectrometer measurements, *Science* **339**, 292 (2013).
- [7] M. Zreda, D. Desilets, T. P. A. Ferré, and R. L. Scott, Measuring soil moisture content non-invasively at intermediate spatial scale using cosmic-ray neutrons, *Geophys. Res. Lett.* **35**, L21402 (2008).
- [8] T. E. Franz, M. Zreda, R. Rosolem, B. K. Hornbuckle, S. L. Irvin, H. Adams, T. E. Kolb, C. Zweck, and W. J. Shuttleworth, Ecosystem-scale measurements of biomass water using cosmic ray neutrons, *Geophys. Res. Lett.* **40**, 3929 (2013).
- [9] M. Andreassen, K. H. Jensen, D. Desilets, T. E. Franz, M. Zreda, H. R. Bogaen, and M. C. Looms, Status and perspectives on the cosmic-ray neutron method for soil moisture estimation and other environmental science applications, *Vadose Zone J.* **16**, 1 (2017).
- [10] R. C. Runkle, A. Bernstein, and P. E. Vanier, Securing special nuclear material: Recent advances in neutron detection and their role in nonproliferation, *J. Appl. Phys.* **108**, 111101 (2010).
- [11] B. M. van der Ende, L. Li, D. Godin, and B. Sur, Stand-off nuclear reactor monitoring with neutron detectors for safeguards and non-proliferation applications, *Nat. Commun.* **10**, 1959 (2019).
- [12] K. Mayer, M. Wallenius, and Z. Varga, Nuclear forensic science: Correlating measurable material parameters to the history of nuclear material, *Chem. Rev.* **113**, 884 (2013).
- [13] J. F. Perello, T. H. Shin, T. Cutler, K. D. Bartlett, C. Bates, D. Beckman, T. C. Borgwardt, D. Hemsing, D. Mercer, and K. Smith, Development of a compact fast-neutron spectrometer for nuclear emergency response applications, *Nucl. Instrum. Methods Phys. Res. A* **1054**, 168468 (2023).
- [14] S. D. Clarke, R. Lopez, V. Mozin, P. Kerr, J. Hutchinson, P. Marleau, and S. A. Pozzi, Neutron spectroscopy of plutonium using a handheld detection system, *Sci. Rep.* **14**, 17192 (2024).

[1] H. Abele, The neutron. Its properties and basic interactions, *Prog. Part. Nucl. Phys.* **60**, 1 (2008).

- [15] M. Wallenius, K. Lützenkirchen, K. Mayer, I. Ray, L. A. de las Heras, M. Betti, O. Cromboom, M. Hild, B. Lynch, A. Nicholl, H. Ottmar, G. Rasmussen, A. Schubert, G. Tamborini, H. Thiele, W. Wagner, C. Walker, and E. Zuleger, Nuclear forensic investigations with a focus on plutonium, *J. Alloys Compd.* **444–445**, 57 (2007), Proceedings of the Plutonium Futures - The Science 2006 Conference.
- [16] H. Meng, Z. Jianyu, W. Jun, and L. Rui, A passive method for the detection of explosives and weapons-grade plutonium in nuclear warheads, *Sci. Glob. Secur.* **26**, 57 (2018).
- [17] E. Lepowsky, J. Jeon, and A. Glaser, Confirming the absence of nuclear warheads via passive gamma-ray measurements, *Nucl. Instrum. Methods Phys. Res. A* **990**, 164983 (2021).
- [18] S. P. LaMont, S. E. Glover, and R. H. Filby, Determination of plutonium-240/239 ratios in low activity samples using high resolution alpha-spectrometry, *J. Radioanal. Nucl. Chem.* **234**, 195 (1998).
- [19] J. Veitch, Parameter estimation for compact binaries with ground-based gravitational-wave observations using the LALInference software library, *Phys. Rev. D* **91**, 042003 (2015).
- [20] G. Ashton, M. Hübner, P. D. Lasky, C. Talbot, K. Ackley, S. Biscoveanu, Q. Chu, A. Divakarla, P. J. Easter, B. Goncharov, F. H. Vivanco, J. Harms, M. E. Lower, G. D. Meadors, D. Melchor, E. Payne, M. D. Pitkin, J. Powell, N. Sarin, R. J. E. Smith, and E. Thrane, Bilby: A user-friendly Bayesian inference library for gravitational-wave astronomy, *ApJS* **241**, 27 (2019).
- [21] R. J. E. Smith, G. Ashton, A. Vajpeyi, and C. Talbot, Massively parallel Bayesian inference for transient gravitational-wave astronomy, *Mon. Not. R. Astron. Soc.* **498**, 4492 (2020).
- [22] R. J. MacDonald and N. Madhusudhan, HD 209458b in new light: Evidence of nitrogen chemistry, patchy clouds and sub-solar water, *Mon. Not. R. Astron. Soc.* **469**, 1979 (2017).
- [23] A. Pinhas, B. V. Rackham, N. Madhusudhan, and D. Apai, Retrieval of planetary and stellar properties in transmission spectroscopy with Aura, *Mon. Not. R. Astron. Soc.* **480**, 5314 (2018).
- [24] R. Trotta, Bayes in the sky: Bayesian inference and model selection in cosmology, *Contemp. Phys.* **49**, 71 (2008).
- [25] U. von Toussaint, Bayesian inference in physics, *Rev. Mod. Phys.* **83**, 943 (2011).
- [26] B. E. Nelson, E. B. Ford, J. Buchner, R. Cloutier, R. F. Díaz, J. P. Faria, N. C. Hara, V. M. Rajpaul, and S. Rukdee, Quantifying the Bayesian evidence for a planet in radial velocity data, *Astron. J.* **159**, 73 (2020).
- [27] G. Ashton *et al.*, Nested sampling for physical scientists, *Nat. Rev. Methods Primers* **2**, 1 (2022).
- [28] M. Praszalowicz, Negative binomial distribution and the multiplicity moments at the LHC, *Phys. Lett. B* **704**, 566 (2011).
- [29] S. V. Tezlař, Significance of the negative binomial distribution in multiplicity phenomena, *Phys. Scr.* **98**, 115310 (2023).
- [30] L. A. Perez, S. Malhotra, J. E. Rhoads, and V. Tilvi, Void probability function of simulated surveys of high-redshift Ly α emitters, *ApJ* **906**, 58 (2021).
- [31] J. N. Fry and S. Colombi, Void statistics and hierarchical scaling in the halo model, *Mon. Not. R. Astron. Soc.* **433**, 581 (2013).
- [32] L. Hurtado-Gil, V. J. Martínez, P. Arnalte-Mur, M.-J. Pons-Bordería, C. Pareja-Flores, and S. Paredes, The best fit for the observed galaxy counts-in-cell distribution function, *A&A* **601**, A40 (2017).
- [33] M. Hameeda, A. Plastino, and M. C. Rocca, Generalized Poisson distributions for systems with two-particle interactions, *IOPSciNotes* **2**, 015003 (2021).
- [34] M. Hünnefeld *et al.*, Combining maximum-likelihood with deep learning for event reconstruction in IceCube, *Proc. Sci.* **395**, 1065 (2022).
- [35] D. Salinas, V. Flunkert, J. Gasthaus, and T. Januschowski, DeepAR: Probabilistic forecasting with autoregressive recurrent networks, *Int. J. Forecast.* **36**, 1181 (2020).
- [36] J. O. Lloyd-Smith, Maximum likelihood estimation of the negative binomial dispersion parameter for highly overdispersed data, with applications to infectious diseases, *PLoS ONE* **2**, e180 (2007).
- [37] E. T. Jaynes, Information theory and statistical mechanics, *Phys. Rev.* **106**, 620 (1957).
- [38] R. Lopez, W. M. Steinberger, N. Giha, P. Marleau, S. D. Clarke, and S. A. Pozzi, Neutron and gamma imaging using an organic glass scintillator handheld dual particle imager, *Nucl. Instrum. Methods Phys. Res. A* **1042**, 167407 (2022).
- [39] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/v6j6-f4rx> for supporting methods and figures, which includes Ref. [40].
- [40] D. Foreman-Mackey, Corner.py: Scatterplot matrices in Python, *J. Open Source Softw.* **1**, 24 (2016).
- [41] D. Servén and C. Brummitt, pyGam: Generalized additive models in Python, Zenodo, 2025, <https://doi.org/10.5281/zenodo.16453423>
- [42] R. Cloutier, N. Astudillo-Defru, R. Doyon, X. Bonfils, J.-M. Almenara, F. Bouchy, X. Delfosse, T. Forveille, C. Lovis, M. Mayor, K. Menou, F. Murgas, F. Pepe, N. C. Santos, S. Udry, and A. Wünsche, Confirmation of the radial velocity super-Earth K2-18c with HARPS and CARMENES, *A&A* **621**, A49 (2019).
- [43] H. Jeffreys, *Theory Of Probability* (Oxford University Press, Oxford, 1948), 2nd ed.
- [44] D. V. Lindley, On a measure of the information provided by an experiment, *Ann. Math. Stat.* **27**, 986 (1956).
- [45] S. P. Luttrell, The use of transinformation in the design of data sampling schemes for inverse problems, *Inverse Probl.* **1**, 199 (1985).
- [46] D. J. C. MacKay, Information-based objective functions for active data selection, *Neural Comput.* **4**, 590 (1992).
- [47] S. Kullback and R. A. Leibler, On information and sufficiency, *Ann. Math. Stat.* **22**, 79 (1951).
- [48] S. A. Pozzi, E. Padovani, and M. Marseguerra, MCNP-PoliMi: A Monte-Carlo code for correlation measurements, *Nucl. Instrum. Methods Phys. Res. A* **513**, 550 (2003).
- [49] J. Allison *et al.*, Recent developments in Geant4, *Nucl. Instrum. Methods Phys. Res. A* **835**, 186 (2016).
- [50] C. Ahdida *et al.*, New capabilities of the FLUKA multi-purpose code, *Front. Phys.* **9**, 788253 (2022).
- [51] D. Breitenmoser, R. Lopez, S. D. Clarke, and S. A. Pozzi, Identifying neutron sources using recoil and time-of-flight

- spectroscopy (Dataset), <https://doi.org/10.5281/zenodo.17905150> (2025).
- [52] J. S. Speagle, Dynesty: A dynamic nested sampling package for estimating Bayesian posteriors and evidences, *Mon. Not. R. Astron. Soc.* **493**, 3132 (2020).
 - [53] E. Higson, W. Handley, M. Hobson, and A. Lasenby, Nestcheck: Diagnostic tests for nested sampling calculations, *Mon. Not. R. Astron. Soc.* **483**, 2044 (2019).
 - [54] J. Skilling, Nested sampling for general Bayesian computation, *Bayesian Anal.* **1**, 833 (2006).
 - [55] E. Higson, W. Handley, M. Hobson, and A. Lasenby, Sampling errors in nested sampling parameter estimation, *Bayesian Anal.* **13**, 873 (2018).
 - [56] D. E. G. Hare, Computing the principal branch of log-gamma, *J. Algorithms* **25**, 221 (1997).
 - [57] P. Virtanen *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nat. Methods* **17**, 261 (2020).
 - [58] R. Lopez, O. Pakari, C. Ballard, S. Clarke, and S. Pozzi, A simulation pipeline for fast neutron imaging and spectroscopy using quantified detector attributes, *Radiat. Meas.* **184**, 107440 (2025).
 - [59] L. Q. Nguyen, G. Gabella, B. L. Goldblum, T. A. Laplace, J. S. Carlson, E. Brubaker, and P. L. Feng, Boron-loaded organic glass scintillators, *Nucl. Instrum. Methods Phys. Res. A* **988**, 164898 (2021).
 - [60] T. H. Shin, P. L. Feng, J. S. Carlson, S. D. Clarke, and S. A. Pozzi, Measured neutron light-output response for *trans*-stilbene and small-molecule organic glass scintillators, *Nucl. Instrum. Methods Phys. Res. A* **939**, 36 (2019).
 - [61] W. K. Warburton, J. S. Carlson, and P. L. Feng, Organic glass scintillator (OGS) property comparisons to stilbene, EJ-276 and BC-404, *Nucl. Instrum. Methods Phys. Res. A* **1018**, 165778 (2021).
 - [62] J. S. Carlson and P. L. Feng, Melt-cast organic glasses as high-efficiency fast neutron scintillators, *Nucl. Instrum. Methods Phys. Res. A* **832**, 152 (2016).
 - [63] M. L. Ruch, P. Marleau, and S. A. Pozzi, Position sensitivity within a bar of stilbene coupled to silicon photomultipliers, in *2016 IEEE Nucl. Sci. Symp. Med. Imaging Conf. Room-Temp. Semicond. Detect. Workshop NSSMICRTSD* (IEEE, Strasbourg, France, 2016), p. 1.
 - [64] N. P. Giha, W. M. Steinberger, L. Q. Nguyen, J. S. Carlson, P. L. Feng, S. D. Clarke, and S. A. Pozzi, Organic glass scintillator bars with dual-ended readout, *Nucl. Instrum. Methods Phys. Res. A* **1014**, 165676 (2021).
 - [65] N. P. Giha, M. L. Ruch, A. D. Fulvio, W. M. Steinberger, and S. A. Pozzi, in *2017 IEEE Nucl. Sci. Symp. Med. Imaging Conf. NSSMIC* (2017), p. 1.
 - [66] W. M. Steinberger, M. L. Ruch, N. Giha, A. D. Fulvio, P. Marleau, S. D. Clarke, and S. A. Pozzi, Imaging special nuclear material using a handheld dual particle imager, *Sci. Rep.* **10**, 1855 (2020).
 - [67] W. Steinberger, N. Giha, M. Hua, S. Clarke, and S. Pozzi, Anisotropic neutron response of *trans*-stilbene and impact on a handheld dual particle imager, *Nucl. Instrum. Methods Phys. Res. A* **1003**, 165266 (2021).
 - [68] F. D. Brooks, R. W. Pringle, and B. L. Funt, Pulse shape discrimination in a plastic scintillator, *IRE Trans. Nucl. Sci.* **7**, 35 (1960).
 - [69] J. K. Polack, M. Flaska, A. Enqvist, C. S. Sosa, C. C. Lawrence, and S. A. Pozzi, An algorithm for charge-integration, pulse-shape discrimination and estimation of neutron/photon misclassification in organic scintillators, *Nucl. Instrum. Methods Phys. Res. A* **795**, 253 (2015).
 - [70] J. B. Birks, Scintillations from organic crystals: Specific fluorescence and relative response to different radiations, *Proc. Phys. Soc. Sect. A* **64**, 874 (1951).