

Deep-learning emulator for photonuclear reactions and a framework for Bayesian joint evaluations via Hamiltonian Monte Carlo

Noah Walton , Ajeeta Khatiwada , and Toshihiko Kawano 

Los Alamos National Laboratory, Bikini Atoll Road, Los Alamos, New Mexico 87544, USA



(Received 10 April 2026; accepted 19 August 2026; published 21 September 2026)

Photonuclear data libraries provide a quantitative description of photon interaction physics for nuclear reaction modeling and underpin a broad range of applications. However, existing evaluations remain limited by persistent inconsistencies between experimental measurements and theoretical descriptions, particularly when multiple isotopes and reaction channels must be treated simultaneously. Correlations induced by shared reaction-model parameters across isotopes, observables, and reaction channels are difficult to rigorously propagate within traditional localized evaluation approaches, motivating joint inference methodologies for physically consistent photonuclear evaluation. In this article, we present a deep-neural network emulator of CoH_3 reaction modeling code that accurately reproduces theoretical calculations with substantial speedup and enables a Markov Chain Monte Carlo inference. We show that a statistically consistent joint Bayesian treatment of ^{235}U and ^{239}Pu photonuclear reaction cross sections, including photofission ratio data, exposes tensions between current nuclear models and the available experimental data. We present a joint photonuclear evaluation that rigorously incorporates these measurements and reports posterior distributions on the underlying reaction-model parameters. Validation studies show high predictive accuracy (root mean square error <0.008 barns) and the ability to generalize to unseen nuclei, indicating that the emulator captures meaningful underlying physics relationships. The joint Bayesian inference reveals that the available experimental data cannot be simultaneously reproduced within the current theoretical description. The inferred posterior distributions identify the model parameters and correlated reaction mechanisms most strongly constrained by the data, exposing tensions across isotopes, observables, and reaction channels that remain hidden in traditional localized evaluations. The analysis also revisits a longstanding ^{235}U photofission cross section discrepancy and reaches conclusions that differ from previous interpretations based on isolated evaluations. More broadly, this work establishes machine-learning-enabled Bayesian inference as a framework for rigorously testing the consistency of photonuclear data and probing the limits of current reaction theory.

DOI: [10.1103/h3f6-zr75](https://doi.org/10.1103/h3f6-zr75)

I. INTRODUCTION

Photonuclear data libraries provide detailed representations of photon interactions with atomic nuclei, leading to particle emission or nuclear disintegration. These interactions play a crucial role in various applications, including fission and fusion energy systems, astrophysics, radiographic techniques, and nuclear safeguards. Despite coordinated international efforts to compile experimental measurements through databases such as EXFOR [1], photonuclear data libraries continue to suffer from sparse and often low-precision experimental coverage. The 2019 update from the Coordinated Research Project (CRP) [2] covers roughly 200 isotopes, with photon energy extending up to 200 MeV. Comparable sublibraries in ENDF/B-VIII.1 [3], TENDL-2025 [4], and JENDL-5 [5] provide additional coverage, although all

remain limited in scope, particularly in covariance representation and coverage at higher energies. Many evaluations remain static carryovers from earlier releases and lack full uncertainty quantification or modern benchmarking. This remains in stark contrast to neutron-induced data libraries, which encompass thousands of isotopes across a much broader energy range and benefit from decades of experimental validation.

Photonuclear reactions are governed by different photon absorption mechanisms whose relative importance depends on the incident photon energy. At photon energies E_γ below approximately 30 MeV, the Giant Dipole Resonance (GDR), a collective excitation in which protons and neutrons oscillate out of phase, dominates. At higher energies, the quasideuteron (QD) process becomes increasingly significant, with the photon interacting with pre-existing proton-neutron pairs within the nucleus, mimicking the photodisintegration of a deuteron [2].

Theoretical modeling of photonuclear processes is typically based on the Compound Nucleus (CN) framework, in which an intermediate nuclear state is formed by the absorption of the incoming photon. This compound nucleus exists in thermal equilibrium, and its subsequent decay is independent

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.

of how the nucleus was formed. The decay is typically described using statistical models such as the Hauser-Feshbach model, analogous to those used in neutron-induced reaction modeling that calculate decay probabilities based on nuclear level densities, particle and photon penetrabilities, energy conservation, and quantum selection rules on spin and parity.

These mechanisms are implemented in reaction codes such as TALYS [6], CoH₃ [7], CCONE [8], and EMPIRE [9], which combine absorption models, pre-equilibrium dynamics, and compound-nucleus decay within a unified framework. Despite the substantial progress, these models remain largely phenomenological and continue to require isotope- and channel-specific parameter tuning. Although certain global parameters exhibit systematic correlations across isotopes, current photonuclear evaluation methodologies still predominantly treat individual nuclei and reaction channels independently. Existing methodologies, largely based on generalized least-squares (GLS) techniques [10,11], rely on simplifying statistical assumptions that are poorly suited to the strongly nonlinear, correlated, and sometimes inconsistent nature of photonuclear data [2,12,13]. As a result, correlated constraints from shared reaction-model physics are not consistently propagated across isotopes, reaction channels, and derived observables such as cross-section ratios.

A deep learning-based emulator offers a pathway to address many of these limitations, not by replacing existing physics-based reaction models, but by serving as a computational surrogate that enables more flexible Bayesian modeling of photonuclear physics and experimental observables. In this paper, we present a deep learning-based emulator trained on detailed outputs from the CoH₃ [7] photonuclear reaction code. The emulator reproduces the underlying physics model with high fidelity while reducing computational cost by more than six orders of magnitude, while retaining full differentiability for gradient-based inference. This computational efficiency enables tractable solutions to complex Bayesian models that consistently connect underlying physics parameters to photonuclear observables across wide energy ranges, multiple reaction channels, multiple isotopes, and nonlinear transformations associated with experimental methods. Within this framework, we leverage Hamiltonian Monte Carlo [14] to estimate the joint Bayesian posterior for physics parameters governing ²³⁹Pu and ²³⁵U (e.g., level densities, fission barriers, and GDR parameters). The framework is not restricted by linear or Gaussian assumptions and naturally accommodates shared and correlated physics parameters within the evaluation. For example, level density parameters for nuclei within the same neutron-emission chain may simultaneously influence multiple neutron multiplicity cross sections for a given target nucleus. Trained across a broad region of the nuclear chart, the emulator reproduces systematic nuclear trends and demonstrates predictive capability for neighboring isotopes outside the training set, indicating that the network learns underlying physical structure rather than merely interpolating between training examples. This behavior suggests that correlated nuclear systematics can be leveraged to inform photonuclear modeling in sparsely constrained regions of the chart of nuclides.

In the study of ²³⁹Pu and ²³⁵U presented herein, no reaction-model parameters are shared directly between isotopes; however, the inclusion of photofission cross-section ratio measurements introduces correlated constraints between the corresponding parameter spaces. The analysis simultaneously incorporates cross sections across six channels $[(\gamma, Xn), (\gamma, 1n), (\gamma, 2n), (\gamma, \text{fis}), (\gamma, \text{Sn}), \text{ and } (\gamma, \text{abs})]$ [2] for both isotopes within a unified Bayesian framework. The resulting posterior distributions quantitatively identify which reaction-model parameters are strongly constrained by the available measurements in these reaction channels and which remain weakly informed. More broadly, the framework enables systematic investigation of correlated constraints, parameter sensitivities, and potential tensions between theoretical descriptions and experimental observables across isotopes and reaction channels. In particular, the longstanding discrepancy in the ²³⁵U (γ, fis) channel is revisited within a framework that incorporates information from multiple reaction channels and isotopes simultaneously.

The emulator architecture and training methodology are described in Sec. II. The Bayesian framework used to condition the photonuclear reaction model on experimental data is presented in Sec. III, followed by results and discussion in Sec. IV, where we examine correlated parameter constraints, model-data tensions, and inconsistencies across reaction channels and isotopes within emulator-driven evaluations. Finally, conclusions are provided in Sec. V.

II. DEEP LEARNING EMULATOR

A. Problem formulation

The deep-learning-based emulator is designed to reproduce the behavior of the CoH₃ photonuclear reaction model with high computational efficiency, enabling consistent application across energies and reaction channels during evaluations. The CoH₃ model predicts photonuclear cross sections for a given nucleus based on a set of nuclear-structure and reaction parameters. Mathematically, the emulator seeks to learn the functional mapping from model inputs X to photonuclear reaction cross sections σ_{ZA} ,

$$X \mapsto \sigma_{ZA}(\gamma, c)(E_\gamma), \quad (1)$$

for all nuclei considered, each uniquely defined by the proton number Z and mass number A . The nuclei considered for this work are defined by an euclidean distance from the line of stability (excluding very light isotopes $A < 20$) and can be visualized later in Fig. 3. For a given nucleus, E_γ is the incident photon energy and c indicates the reaction channel. We define $K = 6$ reaction channels in the set $\mathcal{C}$:

$$\begin{aligned} \mathcal{C} &= \{Xn, 1n, 2n, Sn, \text{abs}, \text{fis}\} \text{ (ordered),} \\ c_k &\in \mathcal{C}, \quad k = 1, \dots, 6. \end{aligned} \quad (2)$$

The physical interpretations of each c_k are as defined in Ref. [2]: Xn is the photoneutron production cross section (including multiplicity), $1n$ and $2n$ are the single and double neutron emission cross sections, Sn is the total photoneutron cross section (not including multiplicity), abs is the total pho-

to absorption cross section, and σ_{fis} is the photofission cross section.

The input parameters, $X_i \in \mathbb{R}^{p_X}$, form a fixed-order vector in which each dimension corresponds to a specific physical model parameter. Consequently, the model is not permutation invariant, as reordering the inputs would alter their physical interpretation and therefore the model output. We can unpack X into smaller fixed-order subsets based on the physics interpretation,

$$X \equiv [\mathcal{N}, \Gamma, \mathcal{F}, \Lambda, \beta], \quad (3)$$

where $\mathcal{N}_i \in \mathbb{R}^{p_N}$ are fundamental nuclear parameters, $\Gamma_i \in \mathbb{R}^{p_\Gamma}$ are GDR parameters, $\mathcal{F}_i \in \mathbb{R}^{p_{\mathcal{F}}}$ are fission-barrier descriptors (if applicable), $\Lambda_i \in \mathbb{R}^{p_\Lambda}$ are level density parameters, and $\beta_i \in \mathbb{R}^{p_\beta}$ are deformation parameters.

The collection of fundamental nuclear parameters, $\mathcal{N}$, is unique to each nucleus and includes known quantities such as the proton and mass numbers and is detailed in the Appendix. In contrast, the other input parameter subsets represent semi-empirical (phenomenological) nuclear model parameters that are not precisely known for any nuclei.

The model input for Γ has a dimension $p_\Gamma = 18$ which describes six GDRs, each with three parameters that characterize the resonance energy, width, and strength. The six GDRs are defined for $E1$, $M1$, $E2$, and $M2$ transitions, with $E1$ and $M1$ transitions represented by up to two Lorentzians (denoted a and b) to account for nuclear deformation [15]. The $E1_a$ and $E1_b$ modes are modeled by a modified Lorentzian, while the remaining modes follow standard Lorentzian parametrizations. The model input for $\mathcal{F}$ has a dimension $p_{\mathcal{F}} = 16$ describing double-hump fission barriers for first, second, third, and fourth chance fission. Each fission barrier is given a height and curvature parameter for each hump, while the inertia parameter is ignored and the default in CoH₃ is taken. Level density information is required not only for the target nucleus but also for the daughter nuclei that appear as intermediate compound systems during particle emission. Channels such as $(\gamma, 1n)$ and $(\gamma, 2n)$ involve sequential decays, and the level densities of each intermediate nucleus affect the branching ratios and final reaction yields. The model input, therefore, takes three level density parameters for each relevant nucleus, resulting in dimension $p_\Lambda = 3$. Finally, there are three deformation parameters (i.e., $\beta = [\beta_2, \beta_4, \beta_6]$) in the input, giving dimension $p_\beta = 3$ corresponding to quadrupole, hexadecapole, and hexacontatetrapole distortions of the nucleus shape [16].

For each nucleus, the prediction target is $y_{ZA} \in \mathbb{R}^q$, which is a vector of cross-section predictions at different photon energies and reaction channels. All reaction channels share a common photon-energy grid of $L = 99$ linearly spaced points from 0.01 to 30 MeV:

$$E_\ell = 0.01 + (\ell - 1)\Delta E, \\ \Delta E = \frac{30 - 0.01}{99 - 1} \approx 0.306 \text{ MeV}, \quad \ell = 1, \dots, 99. \quad (4)$$

For each reaction channel $c \in \mathcal{C}$, let $\sigma_{ZA}(\gamma, c) : \mathbb{R} \rightarrow \mathbb{R}$ denote the cross section (e.g., in barns) for nucleus ZA forming

the channel-by-energy matrix $Y \in \mathbb{R}^{C \times L}$ with entries

$$Y_{k,\ell} = \sigma_{ZA}(\gamma, c_k)(E_\ell), \quad k = 1, \dots, 6, \quad \ell = 1, \dots, 99.$$

The prediction vector $y_{ZA} \in \mathbb{R}^q$ (with $q = K \times L \equiv 594$) is obtained by concatenating the energy-ordered values for each channel (in a fixed channel order):

$$y_{ZA} = [\sigma_{ZA}(\gamma, c_1)(E_1), \dots, \sigma_{ZA}(\gamma, c_1)(E_{99}), \dots, \\ \sigma_{ZA}(\gamma, c_6)(E_1), \dots, \sigma_{ZA}(\gamma, c_6)(E_{99})]^\top = \text{vec}(Y^\top). \quad (5)$$

The full training dataset is then

$$\mathcal{D} = \{(\mathcal{N}_i, \Gamma_i, \mathcal{F}_i, \Lambda_i, \beta_i, Y_i)\}_{i=1}^M, \quad (6)$$

where each entry contains one set of input parameters, including nuclear identifiers $\mathcal{N}_i$, GDR parameters Γ_i , fission-barrier descriptors $\mathcal{F}_i$, level density parameters Λ_i , and deformation parameters β_i , paired with the corresponding CoH₃-generated cross section matrix Y_i . All nuclei (1150) contribute the same number of parameter samples (10 000) to ensure a balanced dataset, for a total dataset size M of 11.5 million.

B. Data preprocessing

Before training the neural network (NN), the dataset $\mathcal{D}$ is scaled and transformed to improve numerical stability and training convergence. The input parameters $(\mathcal{N}_i, \Gamma_i, \mathcal{F}_i, \Lambda_i, \beta_i)$ are normalized to zero mean and unit variance. The prediction targets, $(Y_{i,c}(E))$, where c indices correspond to reaction channels, are similarly first normalized and subsequently transformed to logarithmic space as

$$Y'_{i,c}(E) = \ln \left(1 + \frac{Y_{i,c}(E)}{Y_{0,c}} \right), \quad (7)$$

where $Y_{0,c}$ is a fixed channel-dependent reference value determined from the training set. This transformation compresses the dynamic range of the cross sections and avoids dependence on the absolute scale of $Y_{i,c}$. The transformed targets are then rescaled to the interval $[0, 1]$ for each reaction channel according to

$$\tilde{Y}_{i,c}(E) = \frac{Y'_{i,c}(E) - Y'_{c,\min}}{Y'_{c,\max} - Y'_{c,\min}}. \quad (8)$$

These transformations balance the loss across reaction channels and prevent channels with larger absolute magnitudes from dominating the training. In cases where input parameters are unknown, zero imputation, which corresponds to mean imputation after standardization of the inputs, was applied. The missing features are restricted to quantities such as isospin or J^π that may not be well established for more exotic nuclei. This allows the network to process nuclei with incomplete parameters in subset $\mathcal{N}$ without discarding valid data. In addition, loss masking, where unavailable target elements are excluded from the loss computation, was used to handle cases where elements of the prediction target were missing. Because the network predicts multiple values for each input sample i (i.e., Y_i spans reaction channels k and energy points l), occasional target elements may be unavailable due to computational or numerical errors. Rather than removing the entire

training sample, $\mathcal{D}_i$, the loss contribution from zero-imputed element of $Y_{k,l}$ is set to zero, while the model is allowed to learn from the other elements.

C. Model architecture and training objective

The emulator model was built in PYTHON using the KERAS API [17] with TensorFlow [18] as the backend. The network is structured as a dense (or fully connected) architecture with three hidden layers of widths of 1024, 2048, and 1024 neurons, respectively. Each hidden layer employs a rectified linear unit (ReLU) activation function [19]. The input layer employs batch normalization to stabilize training, while the output layer employs a softplus activation that enforces a physical constraint that all predicted cross sections remain positive.

The loss function combines a mean-squared-error (MSE) term with additional penalty terms that impose soft physical and functional constraints. The model is trained by minimizing the following loss function:

$$\theta^* = \arg \min_{\theta} \frac{1}{M} \sum_{i=1}^M [(\hat{Y}_i - \tilde{Y}_i)^2 + \lambda_{\text{phys}} R_{\text{phys}}(\hat{Y}_i) + \lambda_{\text{dv}} R_{\text{dv}}(\hat{Y}_i, Y_i)], \quad (9)$$

where θ represents the network parameter vector, M is the output dimension, $\hat{Y}$ denotes the network prediction in the transformed target space, and $\tilde{Y}$ is the corresponding prediction in physical cross-section space obtained through the inverse of the transformations given in Eqs. (7) and (8). λ_{phys} and λ_{dv} are non-negative weighting coefficients that control the relative strength of the physics-based and derivative-consistency regularization terms. The physics-informed regularization terms are evaluated in the physical cross section space, $\hat{Y}$, rather than in the transformed training space. The derivative-consistency regularization term is computed as

$$R_{\text{dv}} = \frac{1}{N} \sum_{i=1}^N \left(\frac{\partial \hat{Y}_i}{\partial E_{\gamma}} - \frac{\partial Y_i}{\partial E_{\gamma}} \right)^2, \quad (10)$$

where the derivatives are evaluated using central finite differences for interior points and forward and backward finite differences at the boundaries. The weighting hyperparameter (λ_{dv}) was set to 1.0 such that deviations in function values and derivatives contribute equally to the loss.

The physics penalty R_{phys} calculates the mean-squared violation of the channel constraint,

$$\sigma(\gamma, Xn)(E\gamma) \geq \sigma(\gamma, 1n)(E\gamma) + 2\sigma(\gamma, 2n)(E\gamma) + \sigma(\gamma, \text{fis})(E\gamma). \quad (11)$$

The hyperparameter λ_{phys} was selected by increasing from zero until the constraint was not violated. This was done with a coarse step size because the training was robust against this hyperparameter, and very few violations occurred when $\lambda_{\text{phys}} = 0$.

Explicit regularization, such as L1 and L2 weight penalties and dropout, was not implemented into the loss function or any of the hidden layers as the purpose of this network is to emulate COH₃ exactly rather than to generalize beyond

it, and no symptoms of overfitting were observed in the test set. Early stopping based on a test set, which can sometimes be interpreted as a regularization technique, was employed to reduce computational cost and to enable fair comparison of networks of different sizes during initial NN architecture tuning. While batch normalization and the physics-based soft constraints can have stabilizing effects, they are not employed as regularization; rather, they serve to ensure numerical stability and physical consistency. Optimization was performed using the Adam stochastic gradient descent [19,20] with a learning rate of 0.001.

Training directly on the entire dataset led to slower convergence, so a warm start strategy was adopted to improve the global objective and overall training time. In this approach, θ_0 came from first training on several subsets of $\mathcal{D}$ that each contained only one nuclei (i.e., input parameters $\mathcal{N}$ are all constant and the model learns the relationship for the phenomenological parameters). These models converged to θ_{ZA}^* quickly, and θ_0 was calculated as the average of all θ_{ZA}^* .

D. Generating training data

For a given nucleus, the parameters $\mathcal{N}$ are generally well-known and determined through theoretical and experimental developments that do not involve photoreaction cross sections. These data were compiled from the IAEA Live Chart of Nuclides [21], which is largely informed by the Evaluated Nuclear Structure Data File [22]. Conversely, the parameters Γ , $\mathcal{F}$, Λ , and β are generally determined through, or at least adjusted during, evaluation of experimental photoreaction data. It follows that the goal of the emulator is to find the mapping in Eq. (1) within some reasonable phase space of parameters Γ , $\mathcal{F}$, Λ , and β . For a given nucleus, we define that phase space as a hypercube centered at a vector of prior best estimates for the unknown physics parameters, $\pi_{ZA} = [\Gamma, \mathcal{F}, \Lambda, \beta]$. The size of the phase space is set by a multiplicative parameter u , giving the lower bound $(1 - u)\pi_{ZA}$ and the upper bound $(1 + u)\pi_{ZA}$.

The prior parameter estimates π_{ZA} were determined from a combination of tabulated databases and systematic equations. Because complete databases exist for fission barrier parameters and deformation parameters, central parameter estimates were taken from Refs. [23] and [16], respectively. While the referenced database for fission barrier parameters was intended for neutron-induced reactions, no analogous database exists for photon-induced reactions. The emulator parameter space was therefore selected to be sufficiently broad to encompass parameter values relevant to photon-induced reactions. For an isotope with N neutrons, fission barrier parameters are given for isotopes N through $N - 3$. Isovector GDR parameters ($E1$ transitions) are obtained from a compiled database by the IAEA CRP [24] using experimentally determined parameters when available and systematic equations otherwise. For higher-order GDR parameters, the systematic equations found in Chap. 9 of the TALYS manual [15] were used. If the nucleus is deformed, $E1$ and $M1$ transitions were allowed two GDR modes, both centered at the energies given by the systematic equations. The COH₃ code employs the Gilbert-Cameron level density model [25] by default; the parameters Λ act as scaling




FIG. 1. Schematic of five-fold cross-validation for holdout nuclei. Each column represents onefold with an independent 1/5 split of nuclei used as the held-out test set (blue), while the remaining four splits (gray) form the training data. Within each training set, a stratified 90-10 split of train-validation samples across nuclei is applied to monitor early stopping (illustrated for fold 5).

factors on the asymptotic level density parameter and, for an isotope with N neutrons, are given for isotopes N through $N - 3$. The CoH₃ code automatically recalculates relevant fission level densities when the level density is adjusted.

For this work, we defined the phase space with $u = 0.4$ for all nuclei based on the intuition that the compiled central physics parameters π_{ZA} should be within 40% of the values that best explain measured data (unless in the presence of strong experimental bias). Latin Hypercube Sampling, implemented in SciPy [26], was used to generate samples uniformly across the defined phase space. These samples were then formatted as inputs to CoH₃ to generate labels and target data.

For the CoH₃ calculations, discrete level information up to a default cutoff are taken from the Reference Input Parameter Library (RIPL) database [27]. Calculation of the prompt-fission neutron spectra using the Madland-Nix model [28] was done within CoH₃ to be able to analyze the (γ, Xn) channel because it includes fission multiplicity. Pre-equilibrium particle emission was also included. Other default behavior of the CoH₃ code can be found in its manual and associated publications. The Koning-Delaroch neutron potential [6] was used in these calculations. Other charged particle potentials were neglected in the training data. For high- Z nuclei, such as the actinides evaluated later in this article, this approximation has little-to-no impact. For lighter isotopes, particularly below $Z \approx 30$, this is a poor approximation of reality. Both of these statements are validated in the Appendix. While this physics approximation used to train the emulator limits its applicability to real-world data in the low- Z regime, such nuclei were nevertheless included in the training set to demonstrate the emulator's ability to learn underlying physical trends and exploit correlations across the chart of nuclides.

E. Dataset partitioning strategy

Two separate train-test strategies were used. The first is a 10% test split stratified across nuclei (i.e., each nuclei has 9000 train and 1000 test samples randomly selected). This is a standard strategy, and the resulting model can be seen as the best-performing global emulator. The second strategy employs five-fold cross-validation over held-out nuclei to quantify generalization to previously unseen nuclei. Unlike




FIG. 2. Reaction cross sections predicted by the emulator and calculated with CoH₃. An example of the best, median, and worst cases in the test set is shown (i.e., the 1st, 50th, and 99th quantiles).




FIG. 3. Global emulator test performance (MSE) plotted on the chart of nuclei.

the first strategy, this approach produces five emulator models, each trained on a reduced subset of nuclei; as a result, no single model achieves the same level of global performance as the first strategy. Each of the five folds (split across nuclei) also includes an identical split across samples as in the first strategy: a stratified 10% train-test split across the nuclei in the training set to monitor performance for early stopping (see Fig. 1).

F. Emulator performance

1. Global emulator

This section details the performance of the best-performing global emulator that was built as detailed in the previous sections and trained with the first, standard data partitioning strategy (i.e., no holdout isotopes). To illustrate emulator performance, Fig. 2 compares predicted (emulator) versus true (COH₃) cross sections for test samples at the 1st, 50th, and 99th quantile of mean squared error (MSE). These represent the best-, median-, and worst-performing cases across all nuclei.

More global metrics can provide a broader assessment of emulator performance in different physics regions. For example, Fig. 3 shows the MSE for test samples as a function of the evaluated nucleus. A nearly monotonic trend of worsening performance is seen as nuclei become heavier. Standard error metrics evaluated on the test data are shown in Table I. These metrics are given across the entire test set (aggregate) and for each individual reaction channel in the test set. The metrics are root mean squared error (RMSE), mean absolute error (MAE), median absolute error (MedAE), and mean absolute scaled error (MASE). The MASE provides a scale-invariant metric (with respect to the cross-section magnitudes) that is calculated as the MAE divided by the mean absolute deviation. A relative metric such as mean absolute percent error (MAPE) is not reported because it is not robust against the small values of the reaction cross sections, particularly before thresholds, resulting in a highly skewed metric.

TABLE I. Emulator error metrics evaluated on test data. Metrics are evaluated on the entire test set (aggregate) and for each individual reaction channel.

Channel	RMSE	MAE	MedAE	MASE
Aggregate	0.007 745 8	0.002 138 1	0.000 405 7	0.037 021 7
Xn	0.016 166 2	0.004 537 0	0.001 182 7	0.051 237 3
1n	0.005 177 1	0.002 171 4	0.000 586 2	0.049 083 5
2n	0.004 730 1	0.001 408 3	0.000 169 7	0.070 785 5
Sn	0.004 405 9	0.002 064 7	0.000 667 7	0.036 708 1
Abs	0.004 111 4	0.002 152 5	0.000 891 5	0.034 524 5
Fis	0.003 539 6	0.000 494 4	0.000 035 3	0.089 908 7

The test residuals are shown in Fig. 4 to assess the bias and variance of errors across different reaction channels and incident photon energies. The predictions are mostly unbiased, particularly when assessed using the median residual. The average residual shows small bias near the peaks of several reaction channels, likely skewed by a few large-magnitude residuals that occur in the regions where the reaction cross section itself is large.

A major motivation for the emulator is its dramatic improvement in computational speed. As discussed later in Sec. IV, a standard evaluation of ²³⁹Pu would typically start by generating pointwise cross sections with COH₃ from sampled input parameters (≈ 2500 samples to converge the prior covariance matrix). Table II compares the CPU times for this task using COH₃ and the emulator. Once trained, the emulator offers a speedup of approximately six orders of magnitude.

2. Holdout emulator performance

This section details the performance of the five emulator models, each trained on a smaller subset of nuclei in the five-fold cross-validation holdout study. While none of these emulator models are expected to perform as well as the global emulator, they give a robust indication of the global emulator's ability to learn underlying physics and generalize to unseen nuclei. In each of the following metrics or figures, the performance quantities are computed from the aggregated set of held-out nuclei across all five folds, such that each nucleus appears in the holdout set exactly once, enabling holdout performance to be quantified for all nuclei shown in the chart.

A qualitative visualization of best-, median-, and worst-performing cross-section predictions is given in Fig. 5. Performance across the chart of nuclei is shown for holdout nuclei in Fig. 6, standard error metrics are given in Table III, and residual distributions are shown by reaction channel in Fig. 7. Comparing these characterizations of holdout performance to that of the Global emulator testing, similar trends can be seen [e.g., worsening performance as nuclei become heavier, largest residuals and other error metrics in the (γ , Xn) channel, scale invariant MASE metric shows worst performance in the (γ , fis) channel].




FIG. 4. Characterization of emulator test residuals as a function of incident photon energy for the reaction channels indicated in the panel labels.

III. ENHANCED BAYESIAN EVALUATION

A. Bayesian methods for evaluation

Nuclear data evaluators often take a Bayesian approach, using theory or a previous evaluation as prior knowledge and updating that prior knowledge with experimental data. Analytic and numerical methods exist for solving a Bayesian model, with the former used more prominently in the field of nuclear data evaluation [29]. The analytic approach offers a closed-form solution when the prior and likelihood distributions are in the same conjugate family—often a multivariate normal distribution—and when the relationship between dependent and independent parameters is algebraically linear. This approach is often referred to as Bayesian generalized least squares (GLS). Reference [11] notes two approaches for incorporating a nuclear reaction model into the GLS framework: (1) approximating the model with a first-order Taylor expansion and imposing a prior distribution on the model parameters (this approach is sometimes called generalized linear least squares, GLLS), or (2) sampling model parameters to get the mean and covariance of the corresponding calculated observables, thereby encoding the model constraints and prior knowledge to first order. The former approach pro-

vides posterior estimates for model parameters but must be used in the limit of small parameter perturbations because nuclear reaction models can be highly nonlinear [30]. The latter can be less sensitive to model nonlinearities, but the resulting posterior estimates are on the calculated observable rather than the model parameters. Similarly, the Kalman filter approach [31] is an analytic approach that can offer numerical advantages but often reduces to the same matrix operations used in Bayesian nuclear data evaluation, as in Ref. [32].

Monte Carlo methods, such as importance sampling or Markov chain Monte Carlo (MCMC), provide solutions to Bayesian models that subvert assumptions of Gaussian distributions and linearity. These methods have begun to be explored in nuclear data evaluation for this purpose, particularly for composite observables such as ratio data [33]. MCMC algorithms draw samples directly from the posterior, allowing for efficient sampling of complex posterior distributions with many parameters and/or hierarchical relationships. This permits the inclusion of additional latent variables in the Bayesian evaluation framework to represent deficiencies of the model and/or the data, an approach with only a handful of recent applications in nuclear data evaluation [34,35].

In this work, we leverage the No U-Turn Sampler (NUTS) sampling kernel provided by TENSORFLOW via the PYTHON API [18]. The NUTS kernel dynamically selects hyperparameters for an underlying Hamiltonian Monte Carlo (HMC) sampling kernel, which uses first-order gradient information to take more efficient steps than traditional “random walk” approaches. HMC can provide significantly better convergence for high-dimensional, complex, and/or hierarchical posterior

TABLE II. Comparison of emulator and CoH₃ computation speed.

Model	CPU (s) / 1 sample	CPU (s) / 2500 samples
CoH ₃	76.95	3.85×10^5
Emulator	3.17×10^{-5}	0.16




FIG. 5. Reaction cross sections predicted by the emulator and calculated with COH₃. An example of the best, median, and worst cases in the holdout sets is shown (i.e., the 1st, 50th, and 99th quantiles).

parameter spaces, while NUTS can greatly improve the efficiency and usability of HMC with on-the-fly selection of step size [36,37]. These algorithms require the gradient of the posterior log-probability function for each step, which can be analytic in some cases but is more often done with autodifferentiation in probabilistic programming languages (PPLs) such as STAN [38]. In nuclear data evaluation, however, computing these gradients is often computationally prohibitive or even infeasible because the likelihood depends on high-fidelity nuclear reaction models implemented in external codes that

cannot be readily translated into STAN or other PPLs. A key advantage of the photonuclear reaction cross section emulator is that it provides differentiable surrogate predictions, thereby making gradients readily available and enabling the use of efficient gradient-based sampling algorithms.

B. Demonstration on ²³⁵U and ²³⁹Pu

Available photonuclear reaction measurement data for ²³⁵U and ²³⁹Pu in the EXFOR database [1] were compiled for use in evaluation. A summary of the datasets and their corresponding EXFOR entry numbers is provided in Table IV. It is well recognized in the field of nuclear data evaluation that uncertainties listed in the EXFOR entry can be incomplete or inaccurate. Therefore, we took the following approach: all datasets were assigned their reported statistical uncertainties, while a 10% systematic uncertainty was assigned to those datasets for which no systematic uncertainty was reported. Systematic uncertainties were assumed to be correlated with respect to incident photon energy using a Gaussian correlation model with a damping factor of 0.5, following Ref. [39]. Under this assumption, measurements at nearby




FIG. 6. Holdout emulator performance (MSE) plotted on the chart of nuclei.

TABLE III. Emulator error metrics evaluated on holdout nuclei aggregated over all reaction channels.

RMSE	MAE	MedAE	MASE
0.017 262 1	0.004 264 5	0.000 607 3	0.075 261 6




FIG. 7. Characterization of emulator holdout residuals as a function of incident photon energy for the reaction channels indicated in the panel labels.

photon energies are assigned stronger correlations, with the correlations decreasing smoothly as the energy separation increases. This treatment is necessarily approximate because reported systematic uncertainties are likely attributed to latent observables such as yields or multiplicity counts [2], and conversion of such observables to reaction cross sections involves nontrivial analyses (deconvolution in the case of yield

measurements or solving a linear system of equations in the case of multiplicity counts). Derivative-based approaches to accurately propagate systematic uncertainties require access to the latent observables themselves, which are generally not published. Even when such data are available, the derivatives only enable first-order propagation of uncertainties, which break down when uncertainties are large and analysis

TABLE IV. Experimental data curated for evaluation. Statistical uncertainty and systematic uncertainty columns have an X if the information was provided in the EXFOR entry; otherwise, the attributed uncertainty is reported. The statistical uncertainties for both Katz *et al.* data were found in the literature rather than in the EXFOR entry. Data in the ^{235}U (γ , fis) channel are grayed out because they were excluded from the evaluation; however, they are still visualized in later figures for comparison.

Author	Year	Nuclei	Channel	EXFOR entry	Stat. unc.	Sys. unc.
Katz <i>et al.</i> [40]	1957	^{239}Pu	(γ , Xn)	L0091.005	25%	10%
Katz <i>et al.</i> [40]	1957	^{239}Pu	(γ , fis)	L0085.008	30%	10%
Bowman <i>et al.</i> [41]	1964	^{235}U	(γ , Xn)	M0601.003	X	10%
Bowman <i>et al.</i> [41]	1964	^{235}U	(γ , fis)	M0601.002	X	10%
Shapiro <i>et al.</i> [42]	1971	^{239}Pu	(γ , fis)	L0084.003	X	X
Gurevich <i>et al.</i> [43]	1976	^{235}U	(γ , abs)	M0090.003	X	10%
Gurevich <i>et al.</i> [43]	1976	^{239}Pu	(γ , abs)	M0090.005	X	10%
Caldwell <i>et al.</i> [44]	1980	^{235}U	(γ , Xn)	L0050.009	X	X
Caldwell <i>et al.</i> [44]	1980	^{235}U	(γ , 1n)	L0050.006	X	X
Caldwell <i>et al.</i> [44]	1980	^{235}U	(γ , 2n)	L0050.007	X	X
Caldwell <i>et al.</i> [44]	1980	^{235}U	(γ , Sn)	L0050.032	X	X
Caldwell <i>et al.</i> [44]	1980	^{235}U	(γ , fis)	L0050.008	X	X
Ries <i>et al.</i> [45]	1984	^{235}U	(γ , fis)	M0492.002	X	X
Berman <i>et al.</i> [46]	1986	^{239}Pu	(γ , Xn)	L0058.015	X	X
Berman <i>et al.</i> [46]	1986	^{239}Pu	(γ , 1n)	L0058.012	X	X
Berman <i>et al.</i> [46]	1986	^{239}Pu	(γ , 2n)	L0058.013	X	X
Berman <i>et al.</i> [46]	1986	^{239}Pu	(γ , fis)	L0058.014	X	X
De Moraes <i>et al.</i> [47]	1988	^{239}Pu	(γ , fis)	M0406.003	X	X
Krishichayan <i>et al.</i> [48]	2018	$^{235}\text{U} / ^{239}\text{Pu}$	(γ , fis) ratio	L0233.004	X	X




FIG. 8. Parameters and correlations for the Bayesian posterior with an informed prior distribution on the parameters. The upper figure shows correlations in the posterior parameter distributions, while the lower shows the relative movement of each parameter from its prior value. The gray box in the lower plot shows one relative standard deviation for a given parameter. Definitions and ordering for specific parameters within each parameter subset ($E1$, $M1$, $E2$, $M2$, F , Λ , and β) are given in the Appendix and in Table VI.

involves strongly nonlinear transformations. A more detailed uncertainty analysis is outside the scope of this demonstration but would be valuable if done retroactively for past data and should be considered for any new experimental campaigns.

Experimental data in the ^{235}U (γ , fis) channel were excluded from the evaluation; that is, the data by Bowman (1964), Caldwell (1980), and Ries (1984). Data in this channel were excluded because there are unresolved discrepancies between the different measurements, which have significant implications for the cross-section magnitude where the GDR dominates. This has been a major topic of discussion for evaluators, with modern evaluations like ENDF/B [3,49–51] adopting the reasoning and evaluation of Varlamov [52]. In an attempt to further explore this long-standing issue, we leverage the joint evaluation framework to inform the ^{235}U (γ , fis) channel through other isotopes and channels. By excluding the discrepant data completely, we answer the question: what ^{235}U (γ , fis) cross section is most consistent with an evaluation of other channels?

Lastly, parameter values used to define the center point for the emulator phase space in Sec. IID were taken as the mean of a global Bayesian prior for the physics parameters of each nucleus. All parameters were assigned independent Gaussian priors. GDR parameters were given a standard deviation of 10%, while fission barrier, level density, and deformation parameters were given a standard deviation of 5%. The prior uncertainties assigned to the GDR parameters are approximate estimates informed by the expected accuracy of the systematic relationships used to define their mean values. Similarly, the prior uncertainties for the remaining parameters were chosen based on expert judgment rather than a formal quantitative procedure. Nevertheless, adopting simple prescriptions of 10% and 5% is desirable for reproducibility and transparency. Prior predictive distributions were also used to assess the suitability of these choices. In particular, they indicated that certain parameters, such as the fission barriers, required tighter priors to meaningfully constrain the reaction cross sections. For example, assigning a 10% prior uncertainty




FIG. 9. Posterior-predictive reaction cross section distributions compared with the experimental data for ^{235}U . The prior-predictive is visualized by the gray contours in the background, while the different colored posteriors represent two different treatments in the Bayesian evaluation model (Posterior in blue indicates an informed prior while the MLE in green indicates the maximum likelihood or best fit to the data). The dark red line shows the current ENDF/B VIII.1 [3] evaluation.

to the fission barriers produced an excessively broad prior predictive distribution. This prior π , defined for ^{235}U and ^{239}Pu , was conditioned on the experimental data using a multivariate normal likelihood function with a full covariance given by the statistical and systematic uncertainties defined previously. In addition to this proper posterior model, a maximum likelihood model (maximum likelihood estimator, or MLE) was also fit to the data. That is, the prior parameter model was removed such that the resulting posterior is informed completely by the data likelihood. In this case, the HMC sampler was constrained to the same phase space in which the emulator was trained and validated (i.e., 40% from the centered or mean parameter values) to ensure physically meaningful inference and to avoid the likelihood optimization exploiting unvalidated regions of the emulator input space that could produce artificially good agreement with the data.

IV. EVALUATION RESULTS AND DISCUSSION

The MCMC approach generates samples from the joint posterior across all model parameters. The mean and

standard deviation of parameter priors and posteriors are given in the Appendix, while Fig. 8 shows the correlation between each of the physics parameters in the upper plot, and the lower plot shows the relative movement of the posterior from its prior location for each parameter.

In general, strong positive and negative correlations are observed among the $E1$ GDR parameters. Similarly, the first fission barrier and level density parameters, corresponding to the compound nucleus, with subsequent components along the parameter vector representing daughter nuclei following neutron emission exhibit strong correlations both within and between parameter groups. Many of these intra-isotope correlations have a direct physical interpretation within the underlying reaction model. For example, it makes physical sense that the first and second fission barrier heights show strong correlations; these are positive for ^{235}U and negative for ^{239}Pu , which relates to the observed shape of the data. The nonzero interisotope correlations emerge naturally from the joint posterior conditioned on the fission ratio data and demonstrate that nonlinear experimental constraints propagate




FIG. 10. Posterior-predictive reaction cross section distributions compared with the experimental data for ^{239}Pu . The prior-predictive is visualized by the gray contours in the background, while the different colored posteriors represent two different treatments in the Bayesian evaluation model (Posterior in blue indicates an informed prior while the MLE in green indicates the maximum likelihood or best fit to the data). The dark red line shows the current ENDF/B-VIII.1 [3] evaluation.

consistently across nuclei within the evaluation framework presented here.

The reaction model parameter values in the bottom panel of Fig. 8 show that several parameters are pulled far from their prior, particularly the $E1$ GDR, fission barriers, and level densities. This result suggests some tension between the prior and the data. Their displacement, along with minimal spread in the posterior distribution, provides physically interpretable evidence of which model components are most constrained by the data. On the other hand, we see evidence that $M1$ and higher-order GDR parameters are less constrained by the available reaction data.

Pushing the posterior parameter samples back through the emulator results in a posterior-predictive distribution for reaction cross sections. This is shown in Figs. 9, 10, and 11. The prior-predictive distribution is visualized in gray, showing the 80/20, 90/10, and 95/5 quantiles. These figures show two additional results: (1) the current ENDF/B-VIII.1 evaluation, and (2) the posterior predictive distribution for an evaluation with no prior information; that is, the MLE.

Tension between the prior parameters and the data likelihood was indicated in Fig. 8 by several parameter posteriors moving significantly further than two standard deviations from their prior value. It is similarly seen in the prior predictive quantiles (gray shading in Figs. 9, 10, and 11), as they do not provide good coverage for the data in the $(\gamma, 1n)$, $(\gamma, 2n)$, and (γ, fis) channels for ^{235}U and similarly in the $(\gamma, 1n)$ and $(\gamma, 2n)$ channels for ^{239}Pu . The posterior predictive (blue shading) shows better coverage but still does not fully describe data in the $(\gamma, 1n)$ or $(\gamma, 2n)$ channels for ^{239}Pu or the (γ, Sn) or (γ, fis) channel for ^{235}U . Smaller discrepancies between the posterior predictive and the data are also seen in the bimodality of the (γ, Xn) and (γ, abs) channels for ^{235}U . By removing the prior parameter probability from the Bayesian model, we can solve for the “best possible fit to the data” allowed by the physics model (green shading). This was done by solving the problem with a uninformative prior (very large variance on prior parameter estimates) while also constraining the MCMC samples to the phase space within which the emulator is validated. This constraint is




FIG. 11. Posterior-predictive distribution of the fission ratio (^{235}U to ^{239}Pu) compared with the experimental data used in the evaluation. The prior-predictive is visualized by the gray contours in the background, while the different colored posteriors represent two different treatments in the Bayesian evaluation model (posterior in blue indicates an informed prior while the MLE in green indicates the maximum likelihood or best fit to the data). The dark red line shows the current ENDF/B VIII.1 [3] evaluation.

necessary because, in the absence of a prior on the parameters, maximum-likelihood optimization could identify parameter values outside the validated bounds of the physics emulator that are physically unrealistic and/or provide good fits due to emulator errors rather than true underlying physics.

Even in this case, the resulting posterior predictive estimates fail to reproduce all experimental data within reasonable uncertainty, with a lack of coverage occurring in many of the same channels and energy regions as was seen for the model incorporating an informative prior. In fact, the posterior predictive distributions are nearly identical except for a more pronounced bimodality in the composite reaction channels and an increase in magnitude near the peak of several reaction channels. This similarity indicates that the observed discrepancies are not primarily driven by the prior assumptions but instead reflect limitations intrinsic to the current model description or biases in experimental data. Parameters for the MLE model, shown in Fig. 12, show much broader distributions and often move much further from the prior or centered values. For example, while parameters are still physically possible, the $\pm 40\%$ change seen in several fission barrier heights for ^{235}U is highly unlikely within current theoretical expectations. Collectively, these results indicate that the lack of coverage cannot be attributed to inaccurate prior knowledge alone; rather, they provide evidence of a structural model-data inconsistency.

This is further supported by known deficiencies in the data, including the visually discrepant data in the ^{235}U fission channel. As discussed in Sec. III B, these data were excluded from the evaluation. Still, the resulting reaction cross section informed by all other reaction channels does not agree with any of the available data, nor does it agree with the

ENDF/B evaluation informed by Varlamov in 1987. Our result is largely influenced by correlated constraints from the other reaction channels, particularly the fission ratio data, and suggests that either the available (γ, fis) measurements are inconsistent with the correlated constraints imposed by the remaining experimental data, or the physics model is unable to fully reproduce the underlying physics. The interpretation proposed by Varlamov instead prioritizes agreement with the local (γ, fis) measurements, allowing the evaluation to move further from the underlying reaction-model constraints. In contrast, the present framework enforces simultaneous consistency across correlated observables, reaction channels, and isotopes, substantially reducing the freedom available for localized tuning of the (γ, fis) channel alone. As a result, the preferred (γ, fis) evaluation differs significantly from the Varlamov interpretation. Ultimately, resolving these discrepancies will require either new targeted measurements or further development of the underlying physics model.

Similar indications of model-data inconsistency are also observed in the neutron multiplicity channels. Neutron multiplicity sorting is a well-documented experimental challenge for neutron-emitting channels, where finite neutron-detection efficiency and event misidentification can cause higher-multiplicity reactions to be reconstructed as lower-multiplicity reactions resulting in anticorrelated bias in the $(\gamma, 1n)$ and $(\gamma, 2n)$ channels. This interpretation is supported by the evaluation results, where the largest residual discrepancies and poorest posterior predictive coverage consistently occur in the neutron multiplicity channels after incorporating correlated constraints from the remaining observables. These experimental concerns are known *a priori* but are supported by this investigation because the most egregious indicators of discrepancy in model and data are seen exactly in these neutron multiplicity channels. In particular, the experimental $(\gamma, 1n)$ cross sections systematically exceed the posterior predictive mean while the $(\gamma, 2n)$ cross sections are systematically suppressed, consistent with multiplicity misidentification effects.

To resolve the persistent data discrepancies identified in the partial neutron multiplicity channels and the ^{235}U fission channel, targeted $(\gamma, 1n)$ and $(\gamma, 2n)$ measurements using quasimonoenergetic photon beams and high-efficiency neutron detection are recommended. Where feasible, activation measurements or other redundant observables, such as residual nuclei identification or $n\text{-}\gamma$ coincidence techniques, could be incorporated to provide independent validation of multiplicity assignments. For channels involving fission, advanced charged-particle tracking detectors such as time projection chambers (TPCs), which have demonstrated high-precision fragment identification in neutron-induced fission measurements, could provide improved event characterization and systematic control in photonuclear experiments. Hybrid configurations combining TPC-based fragment detection with high-efficiency neutron arrays would further reduce ambiguity in multiplicity assignments. Finally, experiments delivering raw multiplicity counts, detector timestamps, and response matrices, together with a detailed breakdown of systematic uncertainties would allow for rigorous analysis and propagation of experimental systematics into the evaluation pipeline.




FIG. 12. Parameters and correlations for the maximum likelihood evaluation model. The upper figure shows correlations in the parameter distributions, while the lower shows the relative movement of each parameter from its prior value. The gray box in the lower plot shows one relative standard deviation for a given parameter. Definitions and ordering for specific parameters within each parameter subset ($E1$, $M1$, $E2$, $M2$, F , Λ , and β) are given in the Appendix and in Table VI.

and directly address the multiplicity-sorting ambiguities highlighted here.

The posterior analysis enables a physically meaningful identification of which model parameters require improved validation. Parameters exhibiting large posterior shifts and strong correlations, such as the $E1$ GDR and level density parameters, reflect regions where experimental data exert substantial constraint. In contrast, parameters that undergo large shifts under an uninformative prior yet fail to restore predictive coverage, which include several parameter groups in this analysis, indicate structural limitations of the reaction model or systematic bias in the experimental data, rather than only restrictive prior knowledge. Importantly, many of these parameters govern compound-nucleus physics and are shared with neutron-induced reaction channels. Substantial modification of these quantities would therefore impact established neutron-reaction observables. Nevertheless, despite this expanded parameter freedom, the neutron-multiplicity and fission channels are not fully predictive, indicative of experimental data discrepancies and biases. This distinction

between parameter uncertainty and structural inadequacy is made explicit by the MCMC framework.

V. CONCLUSION

This article detailed the development of deep neural networks for photonuclear reaction cross-section calculations. The global emulator, which can predict cross sections for over one thousand nuclei, is able to emulate COH3 calculations to within a few hundredths of a barn and offers a speedup of six orders of magnitude. Additionally, the network showed a promising ability to generalize to unseen nuclei (i.e., nuclei that were not included in the training set). This result validates that the network is learning underlying physical relationships between input features and reaction cross sections that can be translated from one isotope to another. Furthermore, it points to a promising path for modeling nuclei in regimes where state-of-the-art physics models are unreliable, providing a low-barrier, intermediate-fidelity modeling capability

TABLE V. Ordered parameter sequence and description for fundamental nuclear parameters in subset $\mathcal{N}_f$.

Ind.	Sub. ind.	Parameter	Description
0	0	Z	Proton number of the target nucleus
1	1	A	Mass number of the target nucleus
2	2	BE	Nuclear binding energy of the target nucleus
3	3	ME	Mass excess of the target nucleus
4	4	S_n	Neutron separation energy for target nucleus
5	5	S_p	Proton separation energy for target nucleus
6	6	$S_n (N - 1)$	Neutron separation energy for daughter nucleus after 1-neutron emission from the compound system
7	7	$S_p (N - 1)$	Proton separation energy for daughter nucleus after 1-neutron emission from the compound system
8	8	electric quadrupole	Electric quadrupole moment of the target nucleus
9	9	magnetic dipole	Magnetic dipole moment of the target nucleus
10	10	half-life	Half-life of the target nucleus
11	11	radius	Charge radius of the target nucleus
12	12	J^π	Spin and parity of the ground state of the target nucleus
13	13	isospin	Z component of the isospin of the ground state of the target nucleus

that can be used for exploratory studies and to help guide the development of more rigorous theoretical models.

In this work, the speed and differentiability of the emulator enable more advanced statistical modeling for the inference of key physics parameters and their covariance. We solve a Bayesian model with MCMC, which, if computationally feasible, is preferred to traditional GLS approaches as it provides more detailed information about physics parameters and is not limited by assumptions. Relaxing the assumptions of the GLS approach has several advantages. First, the distribution of photoreaction cross sections is constrained to the physics model exactly; thus, the samples are physical and capture non-Gaussian or nonsymmetric behavior—both of which are expected in the presence of nonlinear physics models. Similarly, ratio data (which are nonlinear with respect to the model by definition) or other observables can confidently be incorporated, and shared physics parameters across isotopes can be easily encoded into the Bayesian model.

The results identify partial neutron-multiplicity channels and ^{235}U photofission as areas of particularly low confidence. Posterior distributions for the physics parameters further reveal which model components are meaningfully constrained by the data, indicating that the M1 and higher-order GDR contributions remain only weakly informed by existing measurements. Our conclusions regarding ^{235}U photofission differ from the interpretation of Varlamov, which has been adopted in the ENDF/B evaluation for many years. While that interpretation prioritizes local agreement with individual (γ , fis) measurements, even when inconsistent with correlated observables and reaction channels, the present analysis instead enforces simultaneous consistency across the coupled multi-channel data and the underlying reaction-model constraints. The broader conclusion is that the currently available data do not uniquely support a single interpretation, emphasizing the need for both new targeted measurements and continued development of photonuclear reaction theory.

Beyond the specific ^{235}U and ^{239}Pu study presented here, the persistence of model-data tension under both informative and uninformative priors highlights the broader importance of explicitly treating model and experimental deficiencies

within Bayesian nuclear data evaluation, rather than implicitly absorbing them through linearized updates. In contrast with first-order GLS approaches, where deviations between theory and experiment are effectively hidden within the update procedure, the MCMC framework enables such discrepancies to be directly identified and quantified. While not explored in detail here, the flexibility of the emulator-MCMC approach also provides a practical pathway for incorporating explicit discrepancy models and for systematically identifying the measurements, reaction channels, and theoretical developments most needed to improve photonuclear reaction physics. Although demonstrated here for photonuclear reactions, the underlying methodology coupling an emulator to an MCMC Bayesian framework is broadly applicable when conditioning a physics model on experimental data, and observed generalization across neighboring isotopes additionally suggests that correlated nuclear systematics learned by the emulator could provide useful constraints in sparsely measured regions of the nuclear chart. Future work will extend these studies for ^{235}U and ^{239}Pu , further develop the modeling framework, and expand the analysis to include ^{238}U , with the goal of more systematically identifying areas requiring additional theoretical development.

ACKNOWLEDGMENTS

Research reported in this presentation was in part supported by the U.S. Department of Energy LDRD program at Los Alamos National Laboratory under Projects No. 20250691ECR and No. 20240878PRD4. This work was carried out under the auspices of the National Nuclear Security Administration of the U.S. Department of Energy at Los Alamos National Laboratory under Contract No. 89233218CNA000001.

DATA AVAILABILITY

Some of the data that support the findings of this article are openly available from the Exfor database. Emulator training and testing data generated with the CoH₃ nuclear reaction

TABLE VI. Prior and posterior model parameters and standard deviations. For the prior, the mean and standard deviation fully describe the Gaussian prior distribution, while the posterior distribution is not strictly Gaussian.

Parameter	U-235		Pu-239	
	Prior	Posterior	Prior	Posterior
<i>E1 Ea</i>	10.8200 ± 1.0820	10.9988 ± 0.0371	11.0700 ± 0.5535	11.1536 ± 0.0262
<i>E1 Ga</i>	3.8800 ± 0.3880	2.7741 ± 0.1077	3.2900 ± 0.3290	2.4790 ± 0.0942
<i>E1 Sa</i>	302.0211 ± 30.2021	299.8238 ± 1.2858	227.7679 ± 22.7768	242.2726 ± 1.3108
<i>E1 Eb</i>	13.8000 ± 1.3800	14.1348 ± 0.0454	14.0000 ± 1.4000	13.9003 ± 0.0408
<i>E1 Gb</i>	4.5400 ± 0.4540	5.1893 ± 0.1166	5.5100 ± 0.5510	5.4363 ± 0.1174
<i>E1 Sb</i>	319.3453 ± 31.9345	333.8649 ± 5.4610	358.1836 ± 35.8184	408.3583 ± 2.3812
<i>M1 Ea</i>	6.6440 ± 0.6644	6.9128 ± 0.5631	6.6067 ± 0.6607	6.0400 ± 0.4387
<i>M1 Ga</i>	4.0000 ± 0.4000	4.1878 ± 0.4394	4.0000 ± 0.4000	3.1941 ± 0.3134
<i>M1 Sa</i>	2.1764 ± 0.2176	2.2342 ± 0.2343	2.2074 ± 0.2207	2.1281 ± 0.2012
<i>M1 Eb</i>	6.6440 ± 0.6644	6.1176 ± 0.5261	6.6067 ± 0.6607	7.2340 ± 0.6387
<i>M1 Gb</i>	4.0000 ± 0.4000	4.9744 ± 0.4873	4.0000 ± 0.4000	3.4755 ± 0.3359
<i>M1 Sb</i>	2.1764 ± 0.2176	2.2491 ± 0.2213	2.2074 ± 0.2207	1.9612 ± 0.2079
<i>E2 Ea</i>	10.2091 ± 1.0209	9.2710 ± 0.8464	10.1518 ± 1.0152	10.0989 ± 0.9891
<i>E2 Ga</i>	3.2900 ± 0.3290	3.2279 ± 0.3126	3.2420 ± 0.3242	3.4327 ± 0.4129
<i>E2 Sa</i>	0.5958 ± 0.0596	0.5887 ± 0.0590	0.6242 ± 0.0624	0.6290 ± 0.0621
<i>M2 Ea</i>	6.6440 ± 0.6644	5.6027 ± 0.6079	6.6067 ± 0.6607	6.8346 ± 0.6131
<i>M2 Ga</i>	4.0000 ± 0.4000	4.0178 ± 0.3632	4.0000 ± 0.4000	4.2209 ± 0.3769
<i>M2 Sa</i>	0.0018 ± 0.0002	0.0017 ± 0.0002	0.0018 ± 0.0002	0.0018 ± 0.0002
<i>Ha (N)</i>	5.7000 ± 0.2850	5.8114 ± 0.0219	5.7500 ± 0.5750	5.7437 ± 0.1571
<i>Ca (N)</i>	0.4000 ± 0.0200	0.4081 ± 0.0209	0.7000 ± 0.0350	0.7304 ± 0.0372
<i>Hb (N)</i>	5.6000 ± 0.2800	4.1793 ± 0.1287	5.5500 ± 0.2775	5.6502 ± 0.0648
<i>Cb (N)</i>	0.3000 ± 0.0150	0.2886 ± 0.0145	0.6000 ± 0.0300	0.6400 ± 0.0317
<i>Ha (N − 1)</i>	6.0500 ± 0.3025	6.0763 ± 0.0373	5.4000 ± 0.2700	4.5824 ± 0.2069
<i>Ca (N − 1)</i>	1.0400 ± 0.0520	1.2504 ± 0.0496	0.7000 ± 0.0350	0.6837 ± 0.0350
<i>Hb (N − 1)</i>	5.4000 ± 0.2700	5.1044 ± 0.2201	5.1000 ± 0.2550	6.4656 ± 0.0477
<i>Cb (N − 1)</i>	0.6000 ± 0.0300	0.6477 ± 0.0328	0.6000 ± 0.0300	0.6084 ± 0.0304
<i>Ha (N − 2)</i>	5.9700 ± 0.2985	4.7874 ± 0.2313	5.8000 ± 0.2900	5.9982 ± 0.2863
<i>Ca (N − 2)</i>	0.8000 ± 0.0400	0.7170 ± 0.0358	0.8000 ± 0.0400	0.8278 ± 0.0418
<i>Hb (N − 2)</i>	5.4500 ± 0.2725	5.7735 ± 0.1870	5.8000 ± 0.2900	5.3859 ± 0.2633
<i>Cb (N − 2)</i>	0.5200 ± 0.0260	0.5164 ± 0.0256	0.5200 ± 0.0260	0.5087 ± 0.0262
<i>Ha (N − 3)</i>	5.8000 ± 0.2900	5.2820 ± 0.2568	6.0000 ± 0.3000	6.5013 ± 0.3268
<i>Ca (N − 3)</i>	1.0400 ± 0.0520	1.1088 ± 0.0568	1.0400 ± 0.0520	1.0781 ± 0.0544
<i>Hb (N − 3)</i>	5.1000 ± 0.2550	5.0252 ± 0.2363	5.0000 ± 0.2500	5.4501 ± 0.2847
<i>Cb (N − 3)</i>	0.6000 ± 0.0300	0.5938 ± 0.0293	0.6000 ± 0.0300	0.6411 ± 0.0327
<i>D (N)</i>	1.0000 ± 0.0500	1.2123 ± 0.0109	1.0000 ± 0.0500	1.1554 ± 0.0133
<i>D (N − 1)</i>	1.0000 ± 0.0500	1.2478 ± 0.0110	1.0000 ± 0.0500	1.1553 ± 0.0146
<i>D (N − 2)</i>	1.0000 ± 0.0500	1.0171 ± 0.0348	1.0000 ± 0.0500	0.8719 ± 0.0348
<i>D (N − 3)</i>	1.0000 ± 0.0500	0.9906 ± 0.0487	1.0000 ± 0.0500	0.9579 ± 0.0495
<i>beta0</i>	0.2141 ± 0.0107	0.2098 ± 0.0105	0.2223 ± 0.0111	0.2093 ± 0.0108
<i>beta1</i>	0.1092 ± 0.0055	0.1019 ± 0.0050	0.0949 ± 0.0047	0.1007 ± 0.0050
<i>beta2</i>	−0.0048 ± −0.0002	−0.0054 ± 0.0003	−0.0180 ± −0.0009	−0.0182 ± 0.0009

modeling code are not publicly available upon publication because it is not technically feasible and/or the cost of preparing, depositing, and hosting the data would be prohibitive within the terms of this research project. The data are available from the authors upon reasonable request.

APPENDIX

Section II detailed the development of a deep neural network emulator for photonuclear reaction cross sections that spans multiple nuclei and a significant phase space of input

parameters. The network's input features, X , are an ordered sequence that can be separated into ordered subsets as

$$X \equiv [\mathcal{N}, \Gamma, \mathcal{F}, \Lambda, \beta]. \quad (\text{A1})$$

Table V describes the ordered sequence for the fundamental nuclear parameters used by the emulator. Prior and posterior parameter values, along with uncertainties, are shown in Table VI. In this table, the *E1*, *M1*, *E2*, and *M2* GDR energies (in MeV) are given by *Ea* (or *Eb* if a bimodal GDR is included), the widths (in MeV) are given by *Ga* (or *Gb* if bimodal), and the peak value of the Lorentzian (in mb) is




FIG. 13. External validation of Emulator predictions for ^{235}U posterior parameter set. Posterior parameters inferred using the Emulator or MCMC framework were used in CoH₃ to calculate photonuclear reaction cross sections (black) and validate the emulator predictions (green). Additionally, the CoH₃ calculation was done including charged-particle potentials (red) to show negligible difference in the actinide region.

given by Sa (or Sb if bimodal). These parameters are ordered as given in Table VI and belong to the network feature subset Γ in Eq. (A1). The double-humped fission barrier parameters are the heights (Ha or Hb) and curvatures (Ca or Cb), both in MeV, given for nuclei (N) to ($N - 3$), denoted in parentheses. Similarly, the level density scaling factors D (unitless) are given for nuclei (N) to ($N - 3$), as denoted in the parentheses. Lastly, the static deformation parameters for the target nucleus are given by the beta parameters. The fission barrier parameters, level density scaling factors, and static deformation parameters are all ordered as given in Table VI and belong to the network feature subsets in Eq. (A1), $\mathcal{F}$, Λ , and β , respectively.

The posterior mean values were input to CoH₃ as a means of further validating the emulator performance and the methodology as a whole. The results in Figs. 13 and 14 compare the emulator and CoH₃ predictions at these parameter values. Good performance is seen, similar to the performance in the test set. Finally, it was noted that the global emulator was trained and tested on nuclei across the chart; however, the data generated with CoH₃ did not include charged-particle

potentials. For real-world applications in the high- Z and actinide regions, this should be negligible at the incident photon energy domain of interest due to the Coulomb barrier. This is validated in Figs. 13 and 14 because an additional calculation was done including the following particle potentials in CoH₃:

- (1) Proton: Koning
- (2) Alpha: Avrigeanu2009
- (3) Deuteron: Bojowald
- (4) Helium-3: Becchetti
- (5) Triton: Becchetti

For ^{235}U and ^{239}Pu , these additional particle potentials have no impact on the calculated photoreaction cross sections. These are expected to have an impact for lower- Z nuclei and would limit the applicability of this specific emulator for real-world evaluation. It is important to note that the data generated from the physics model did not include charged particle emission; therefore, the emulator does not learn these effects. However, it is not tested on those effects either since the test data was partitioned from the same set of physics model calculations. Figure 15 shows a comparison




FIG. 14. External validation of Emulator predictions for ^{239}Pu posterior parameter set. Posterior parameters inferred using the Emulator or MCMC framework were used in CoH_3 to calculate photonuclear reaction cross sections (black) and validate the emulator predictions (green). Additionally, the CoH_3 calculation was done including charged particle potentials (red) to show negligible difference in the actinide region.

of CoH_3 calculations at the default or centered parameters, both with and without the charged particle potentials for a few nuclei at lower Z (^{39}K , ^{63}Cu , and ^{74}Ge). Inclusion of these charged-particle potentials reduces the neutron production

cross sections for ^{39}K and ^{63}Cu but becomes negligible for ^{74}Ge . Of course, the photoabsorption cross section is the same regardless.




FIG. 15. Comparison of the impact of excluding charge particle potentials for nuclei with different Z . Solid lines include charged-particle potentials, while dotted lines with black borders indicate reaction calculations using only the neutron potential. The colors indicate different nuclei.

- [1] N. Otuka, E. Dupont, V. Semkova, B. Pritychenko, A. I. Blokhin, M. Aikawa, S. Babykina, M. Bossant, G. Chen, S. Dunaeva, R. A. Forrest, T. Fukahori, N. Furutachi, S. Ganesan, Z. Ge, O. O. Gritzay, M. Herman, S. Hlavač, K. Kato, B. Lalremruata, *et al.*, Towards a more complete and accurate experimental nuclear reaction data library (EXFOR): International collaboration between nuclear reaction data centres (NRDC), *Nucl. Data Sheets* **120**, 272 (2014).
- [2] T. Kawano, Y. S. Cho, P. Dimitriou, D. Filipescu, N. Iwamoto, V. Plujko, X. Tao, H. Utsunomiya, V. Varlamov, R. Xu, R. Capote, I. Gheorghe, O. Gorbachenko, Y. L. Jin, T. Renstrøm, M. Sin, K. Stopani, Y. Tian, G. M. Tveten, J. M. Wang, *et al.*, IAEA photonuclear data library 2019, *Nucl. Data Sheets* **163**, 109 (2020).
- [3] G. P. A. Nobre, R. Capote, M. T. Pigni, A. Trkov, C. M. Mattoon, D. Neudecker, D. A. Brown, M. B. Chadwick, A. C. Kahler, N. A. Kleedtke, M. Zerkle, A. I. Hawari, C. W. Chapman, N. C. Fleming, J. L. Wormald, K. Ramić, Y. Danon, N. A. Gibson, P. Brain, M. W. Paris, *et al.*, ENDF/B-VIII.1: Updated nuclear reaction data library for science and applications, *Nucl. Data Sheets* **210**, 1 (2026).
- [4] A. J. Koning, D. Rochman, J.-C. Sublet, N. Dzysiuk, M. Fleming, and S. van der Marck, TENDL: Complete nuclear data library for innovative nuclear science and technology, *Nucl. Data Sheets* **155**, 1 (2019).
- [5] O. Iwamoto, N. Iwamoto, S. Kunieda, F. Minato, S. Nakayama, Y. Abe, K. Tsubakihara, S. Okumura, C. Ishizuka, T. Yoshida, S. Chiba, N. Otuka, J.-C. Sublet, H. Iwamoto, K. Yamamoto, Y. Nagaya, K. Tada, C. Konno, N. Matsuda, K. Yokoyama, *et al.*, Japanese evaluated nuclear data library version 5: JENDL-5, *J. Nucl. Sci. Technol. (Abingdon, U. K.)* **60**, 1 (2023).
- [6] A. J. Koning and J. P. Delaroche, Local and global nucleon optical models from 1 keV to 200 MeV, *Nucl. Phys. A* **713**, 231 (2003).
- [7] T. Kawano, *CoH₃: The coupled-channels and Hauser-Feshbach code*, *Compound-Nuclear Reactions*, Springer Proceedings in Physics, Vol. 310 (Springer, Cham, 2021), pp. 27–34.
- [8] O. Iwamoto, N. Iwamoto, S. Kunieda, F. Minato, and K. Shibata, The CCONE code system and its application to nuclear data evaluation for fission and other reactions, *Nucl. Data Sheets* **131**, 259 (2016).

- [9] M. Herman, R. Capote, B. V. Carlson, P. Obložinský, M. Sin, A. Trkov, H. Wienke, and V. Zerkin, EMPIRE: Nuclear reaction model code system for data evaluation, *Nucl. Data Sheets* **108**, 2655 (2007).
- [10] B. Faleschini and M. Petilli, MARA: A generalized least squares method for nuclear data adjustment, Tech. Rep. No. RT/FI(73)33 (Comitato Nazionale per l'Energia Nucleare, Rome, Italy, 1973).
- [11] G. Schnabel and H. Leeb, A modified generalized least squares method for large scale nuclear data evaluation, *Nucl. Instrum. Methods Phys. Res., Sect. A* **841**, 87 (2017).
- [12] V. V. Varlamov, A. I. Davydov, and V. N. Orlin, Photonuclear partial reaction cross sections: Systematic uncertainties and reliability, *At. Data Nucl. Data Tables* **161**, 101697 (2025).
- [13] T. Kawano, Evaluation of photo-induced reactions on $^{39,41}\text{K}$, Tech. Rep. No. LA-UR-14-20531 (Los Alamos National Laboratory, Los Alamos, 2014).
- [14] M. Betancourt, A conceptual introduction to Hamiltonian Monte Carlo, [arXiv:1701.02434](https://arxiv.org/abs/1701.02434).
- [15] A. J. Koning, S. Hilaire, and S. Goriely, TALYS: Modeling of nuclear reactions, *Eur. Phys. J. A* **59**, 131 (2023).
- [16] P. Moller, J. R. Nix, W. D. Myers, and W. J. Swiatecki, Nuclear ground-state masses and deformations, *At. Data Nucl. Data Tables* **59**, 185 (1995).
- [17] F. Chollet *et al.*, KERAS, <https://keras.io> (2015).
- [18] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, *et al.*, TensorFlow: Large-scale machine learning on heterogeneous systems, <https://www.tensorflow.org/> (2015), software available from tensorflow.org.
- [19] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, 2016).
- [20] D. P. Kingma and J. Ba, Adam: A method for stochastic optimization, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [21] International Atomic Energy Agency, Live chart of nuclides, <https://www-nds.iaea.org/relnsd/vcharthtml/VChartHTML.html> (2026), accessed 2026/09/03 01:12:34.
- [22] National Nuclear Data Center, Evaluated nuclear structure data file (ENSDF) (2022).
- [23] O. Iwamoto, T. Nakagawa, N. Otuka, S. Chiba, K. Okumura, G. Chiba, T. Ohsawa, and K. Furutaka, JENDL actinoid file 2008, *J. Nucl. Sci. Technol. (Abingdon, U. K.)* **46**, 510 (2009).
- [24] V. A. Plujko, O. M. Gorbachenko, R. Capote, and P. Dimitriou, Giant dipole resonance parameters of ground-state photoabsorption: Experimental values with uncertainties, *At. Data Nucl. Data Tables* **123-124**, 1 (2018).
- [25] A. Gilbert and A. G. W. Cameron, A composite nuclear-level density formula with shell corrections, *Can. J. Phys.* **43**, 1446 (1965).
- [26] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, *et al.*, SciPy 1.0: Fundamental algorithms for scientific computing in PYTHON, *Nat. Methods* **17**, 261 (2020).
- [27] R. Capote, M. Herman, P. Obložinský, P. G. Young, S. Goriely, T. Belgia, A. V. Ignatyuk, A. J. Koning, S. Hilaire, V. A. Plujko, M. Avrigeanu, O. Bersillon, M. B. Chadwick, T. Fukahori, Z. Ge, Y. Han, S. Kailas, J. Kopecky, V. M. Maslov, G. Reffo, *et al.*, RIPL – Reference input parameter library for calculation of nuclear reactions and nuclear data evaluations, *Nucl. Data Sheets* **110**, 3107 (2009).
- [28] D. G. Madland and J. R. Nix, New calculation of prompt fission neutron spectra and average prompt neutron multiplicities, *Nucl. Sci. Eng. (La Grange Park, IL, U. S.)* **81**, 213 (1982).
- [29] R. Capote, D. L. Smith, and A. Trkov, Nuclear data evaluation methodology including estimates of covariances, *EPJ Web Conf.* **8**, 04001 (2010).
- [30] N. M. Larson, Updated users' guide for SAMMY: Multilevel R-matrix fits to neutron data using Bayes' equations, Tech. Rep. No. ORNL/TM-9179/R8 (Oak Ridge National Laboratory, Oak Ridge, 2008).
- [31] R. E. Kalman, A new approach to linear filtering and prediction problems, *J. Basic Eng.* **82**, 35 (1960).
- [32] A. E. Lovell, T. Kawano, P. Talou, and G. Y. Rusev, Fission product yield modeling and evaluation, *EPJ Web Conf.* **322**, 07003 (2025).
- [33] R. Capote, D. L. Smith, A. Trkov, and M. Meghizifene, A new formulation of the unified Monte Carlo approach (UMC-B) and cross-section evaluation for the dosimetry reaction $^{55}\text{Mn}(n, \gamma)^{56}\text{Mn}$, *J. ASTM Int.* **9**, 1 (2012).
- [34] J. M. Brown, G. Arbanas, A. Holcomb, and D. Wiarda, Bayesian Monte Carlo evaluation framework for imperfect nuclear data, *Trans. Am. Nucl. Soc.* **125**, 686 (2021).
- [35] N. A. W. Walton, D. Neudecker, S. A. V. Wiel, M. J. Grosskopf, and K. J. Kelly, Machine learning-assisted identification of potential sources of bias in measurements of prompt-fission neutron spectra, *Comput. Phys. Commun.* **315**, 109698 (2025).
- [36] M. D. Hoffman and A. Gelman, The No-U-turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.* **15**, 1593 (2014).
- [37] M. Betancourt, A conceptual introduction to Hamiltonian Monte Carlo, [arXiv:1701.02434](https://arxiv.org/abs/1701.02434).
- [38] B. Carpenter, A. Gelman, M. D. Hoffman, D. Lee, B. Goodrich, M. Betancourt, M. Brubaker, J. Guo, P. Li, and A. Riddell, Stan: A probabilistic programming language, *J. Stat. Soft.* **76**, 1 (2017).
- [39] D. Neudecker, ARIADNE – A program estimating covariances in detail for neutron experiments, *EPJ Nuclear Sci. Technol.* **4**, 34 (2018).
- [40] L. Katz, K. G. McNeill, M. LeBlanc, and F. Brown, Photoneutron emission from ^{232}Th , ^{233}U , ^{238}U , and ^{239}Pu , *Can. J. Phys.* **35**, 470 (1957).
- [41] C. D. Bowman, G. F. Auchampaugh, and S. C. Fultz, Photodisintegration of ^{235}U , *Phys. Rev.* **133**, B676 (1964).
- [42] A. Shapiro and W. F. Stubbs, Photofission cross section of plutonium-238 and plutonium-239, *Nucl. Sci. Eng. (La Grange Park, IL, U. S.)* **45**, 47 (1971).
- [43] G. Gurevich, L. Lazareva, V. Mazur, G. Solodukhov, and B. Tulupov, Giant resonance in the total photoabsorption cross section of $Z \approx 90$ nuclei, *Nucl. Phys. A* **273**, 326 (1976).
- [44] J. T. Caldwell, E. J. Dowdy, B. L. Berman, R. A. Alvarez, and P. Meyer, Giant resonance for the actinide nuclei: Photoneutron and photofission cross sections for ^{235}U , ^{236}U , ^{238}U , and ^{232}Th , *Phys. Rev. C* **21**, 1215 (1980).
- [45] H. Ries, U. Kneissl, G. Mank, H. Ströher, W. Wilke, R. Bergère, P. Bourgeois, P. Carlos, J. Fallou, P. Garganne, A. Veyssière, and L. Cardman, Absolute photofission cross sections of $^{235,238}\text{U}$ measured with tagged photons between 40 and 105 MeV, *Phys. Lett. B* **139**, 254 (1984).

- [46] B. L. Berman, J. T. Caldwell, E. J. Dowdy, S. S. Dietrich, P. Meyer, and R. A. Alvarez, Photofission and photoneutron cross sections and photofission neutron multiplicities for ^{233}U , ^{234}U , ^{237}Np , and ^{239}Pu , *Phys. Rev. C* **34**, 2201 (1986).
- [47] M. A. P. De Moraes and M. T. Cesar, Photofission cross sections of ^{233}U and ^{239}Pu near threshold induced by gamma rays from thermal neutron capture, *Nucl. Instrum. Methods Phys. Res., Sect. A* **277**, 467 (1989).
- [48] Krishichayan, S. W. Finch, C. R. Howell, A. P. Tonchev, and W. Tornow, Monoenergetic photon-induced fission cross-section ratio measurements for ^{235}U , ^{238}U , and ^{239}Pu from 9.0 to 17.0 MeV, *Phys. Rev. C* **98**, 014608 (2018).
- [49] D. A. Brown, M. B. Chadwick, R. Capote, A. C. Kahler, A. Trkov, M. W. Herman, A. A. Sonzogni, Y. Danon, A. D. Carlson, M. Dunn, D. L. Smith, G. M. Hale, G. Arbanas, R. Arcilla, C. R. Bates, B. Beck, B. Becker, F. Brown, R. J. Casperson, J. Conlin, *et al.*, ENDF/B-VIII.0: The 8th major release of the nuclear reaction data library with CIELO-project cross sections, new standards and thermal scattering data, *Nucl. Data Sheets* **148**, 1 (2018).
- [50] M. B. Chadwick, P. Obložinský, M. Herman, N. M. Greene, R. D. McKnight, D. L. Smith, P. G. Young, R. E. MacFarlane, G. M. Hale, S. C. Frankle, A. C. Kahler, T. Kawano, R. C. Little, D. G. Madland, P. Moller, R. D. Mosteller, P. R. Page, P. Talou, H. Trellue, M. C. White, *et al.*, ENDF/B-VII.0: Next generation evaluated nuclear data library for nuclear science and technology, *Nucl. Data Sheets* **107**, 2931 (2006).
- [51] M. B. Chadwick, M. Herman, P. Obložinský, M. E. Dunn, Y. Danon, A. C. Kahler, D. L. Smith, B. Pritychenko, G. Arbanas, R. Arcilla, R. Brewer, D. A. Brown, R. Capote, A. D. Carlson, Y. S. Cho, H. Derrien, K. Guber, G. M. Hale, S. Hoblit, S. Holloway, *et al.*, ENDF/B-VII.1 nuclear data for science and technology: Cross sections, covariances, fission product yields and decay data, *Nucl. Data Sheets* **112**, 2887 (2011).
- [52] V. Varlamov, N. G. Efimkin, V. V. Surgutanov, A. A. Khoronenko, and A. P. Chernyaev, Photonuclear data: Photofission of $^{235,238}\text{U}$, Tech. Rep. No. CCDFE/FIS2 (Center for Photonuclear Experiments Data, Moscow State University, Moscow, Russia, 1987).