

<https://doi.org/10.1038/s42005-026-02499-6>

Simulation-based inference for precision neutrino physics through neural Monte Carlo tuning

Check for updates

Arsenii Gavrikov¹✉, Andrea Serafini¹✉, Dmitry Dolzhikov^{2,3}✉, Alberto Garfagnini¹, Maxim Gonchar^{2,3}, Marco Grassi¹ & JUNO-Italia Consortium*

Precise modeling of detector energy response is crucial for next-generation neutrino experiments, which present computational challenges due to the lack of analytical likelihoods. We propose a solution using neural likelihood estimation within the simulation-based inference framework. We develop two complementary neural density estimators that model likelihoods of calibration data: conditional normalizing flows and a transformer-based regressor. We adopt JUNO — a large neutrino experiment — as a case study. The energy response of JUNO depends on several parameters, all of which should be tuned, given their non-linear behavior and strong correlations in the calibration data. To this end, we integrate the modeled likelihoods with Bayesian nested sampling for parameter inference, achieving uncertainties limited only by statistics with near-zero systematic biases. The normalizing flows model enables unbinned likelihood analysis, while the transformer provides an efficient binned alternative. By providing both options, our framework offers flexibility to choose the most appropriate method for specific needs. Finally, our approach establishes a template for similar applications across experimental neutrino and broader particle physics.

The new precision era in neutrino physics¹ demands accurate detector energy response characterization. The relationship between the energy released by a neutrino interaction in the detector (deposited energy) and the measured experimental quantities rarely follows a simple analytical form in high-precision detectors, instead representing a complex function, dependent on detector geometry, material properties, and interacting particles.

The absence of analytical descriptions necessitates sophisticated Monte Carlo (MC) simulations that account for a wide range of physics effects. However, these simulations may depend on multiple parameters that summarize complex behaviors. Despite implementing detector properties based on the best available knowledge, inherent uncertainties and necessary simplifications remain. These limitations mean that simulations typically require further refinement to accurately reproduce experimental observations through a process known as MC tuning, which involves systematic adjustment of simulation parameters to achieve agreement between simulated outputs and experimental data.

For MC tuning, parameter inference requires a likelihood function $p(x|\phi)$, where x represents either simulated or observed data and ϕ are the parameters to be tuned. Once constructed, this likelihood enables recovery of the parameter values ϕ that best match the observed data x . However, this

likelihood function is computationally intractable directly from simulations. Approaches employed in neutrino experiments of the previous generation, like Borexino² and Daya Bay³, involve iterative MC tuning where parameters are adjusted manually to match calibration data. However, this becomes computationally prohibitive for modern experiments.

Simulation-Based Inference (SBI, or likelihood-free inference)^{4–9} addresses this challenge by learning likelihood or posterior approximations directly from simulations. These techniques are now widely used across multiple scientific domains, including cosmology^{10–12}, astrophysics^{13–15}, accelerator^{16–19}, and neutrino physics^{20,21}, and beyond^{22,23}. Within SBI, Neural Likelihood Estimation (NLE) constructs approximations to the likelihood function $p(x|\phi)$, while Neural Posterior Estimation (NPE) learns the posterior $q(\phi|x)$ directly and is attractive for ultra-fast inference. NLE separates likelihood estimation from statistical inference, retaining full control over statistical interpretation and uncertainty quantification. The choice between NPE and NLE is guided by data dimensionality, simulation cost, and the need for transparent statistical treatment.

The Jiangmen Underground Neutrino Observatory (JUNO)^{24,25} exemplifies these challenges. The experiment consists of a central detector containing a 2×10^7 kg liquid scintillator-based (LS) target within a 35.4 m

¹INFN, Sezione di Padova e Università di Padova, Dipartimento di Fisica e Astronomia, Padua, Italy. ²Joint Institute for Nuclear Research, Dubna, Russia. ³Institute for Nuclear Research of the Russian Academy of Sciences, Moscow, Russia. *A list of authors and their affiliations appears at the end of the paper.

✉ e-mail: arsenii.gavrikov@pd.infn.it; andrea.serafini@pd.infn.it; ddolzhikov@jinr.ru

diameter acrylic sphere, monitored by 17,612 20" and 25,600 3" photo-multiplier tubes (PMTs), converting light produced in the detector medium to electric charge. This design yields exceptional light collection, quantified by the total number of photons converted to electrons by the PMTs, i.e., the number of photo-electrons $N_{p.e.}$. In our analysis, we exploit only the 20" PMTs, which provide the most precise energy measurement by detecting, on average, 1600 photo-electrons when a 1 MeV gamma loses all its energy at the center of the detector. However, incorporating the 3" PMTs is straightforward if needed. JUNO's primary goals include determining the neutrino mass ordering²⁶ and measuring three oscillation parameters (Δm_{21}^2 , Δm_{31}^2 , $\sin^2\theta_{12}$) with sub-percent precision²⁷, alongside a broad program exploring phenomena from geoneutrinos to supernovae^{28–35}. As a large-scale, next-generation LS detector, it requires exceptional energy resolution and control over energy-related systematic uncertainties²⁵ to achieve its goals. Thus, JUNO's complex, non-uniform, and non-linear response³⁶, coupled with its strong reliance on highly precise MC simulations, makes it a suitable platform for developing and validating SBI approaches.

The energy response in JUNO is characterized by three key parameters with different physical origins and effects: (1) the Birks' coefficient (k_B), which models the non-linear energy quenching at high ionization densities, where part of the energy deposited by the particle is dissipated as heat in the LS rather than being converted into scintillation light; (2) the LS absolute light yield (Y), which defines the number of scintillation photons emitted per unit energy after quenching; (3) the Cherenkov light yield factor (f_C), which scales the energy-dependent yield of photons originating from Cherenkov radiation. Both k_B and f_C model the non-linear relation between the deposited energy and the collected charge, while Y scales energy response linearly³⁶. Precisely calibrating these parameters within JUNO's MC simulation framework³⁷ is pivotal for mitigating systematic uncertainties. Furthermore, accurate simulations are important for any machine learning (ML) applications used in JUNO^{38–44}, preventing performance degradation due to discrepancies between simulated training data and real experimental data.

To constrain the energy response model and determine the optimal (k_B, f_C, Y) parameters, JUNO employs a comprehensive calibration program using radioactive sources deployed at various positions within the detector⁴⁵. The strong correlations between k_B, f_C , and Y require a combined analysis of multiple sources with different energy spectra to break parameter degeneracies. Further details can be found in the JUNO's calibration strategy paper⁴⁵.

For each event of the calibration data, the total signal collected by the PMT array can be summarized by the total number of photo-electrons $N_{p.e.}$ detected by all channels. Since the energy response is modeled at the event level, this compression of the full PMT-wise information into one-dimensional $N_{p.e.}$ distributions enables efficient likelihood modeling while preserving the information relevant for parameter estimation. $N_{p.e.}$ serves as a proxy for the visible energy of an event and translates to the x in the notation used above, while ϕ is the vector of the energy response parameters (k_B, f_C, Y).

In this work, we employ the NLE method to address the challenge of MC tuning. We first train a model for the likelihood estimation using simulation data and then combine that model with conventional statistical tools to obtain the posterior. We use simulated data corresponding to five different calibration sources placed at the detector center: three gamma sources (¹³⁷Cs, ⁴⁰K, ⁶⁰Co) and two neutron sources (²⁴¹Am-Be, ²⁴¹Am-¹³C). For neutron sources, the final state usually consists of an excited state (¹²C*, ¹⁶O*, respectively) emitting a gamma, and a neutron, which is captured by hydrogen (n-H) or carbon (n-¹²C) atoms in the LS, leading to the emission of an additional gamma. Note that the light emission properties of the scintillator are position-independent, and the detector center is chosen to follow JUNO's calibration strategy, where some sources are deployed exclusively at the center. Should position-dependent calibration be desired to further increase dataset statistics, our method naturally extends by treating source position as an additional conditioning variable. We develop two

complementary neural density estimators that construct approximations to the $p(x|\phi)$ function using calibration data: a Transformer Encoder Density Estimator (TEDE)⁴⁶ that maps energy response parameters to histogram-based density approximations, and a Normalizing Flows Density Estimator (NFDE)^{47–49} that models continuous probability densities for unbinned analysis. We integrate these learned likelihood approximations with Bayesian nested sampling^{50,51} for inference of the energy response parameters. This paper also represents a substantial extension of our previous work⁵², where we studied TEDE and a Wasserstein generative adversarial network with gradient penalty^{53,54}.

Both developed NLE methods successfully recover JUNO's energy response parameters (k_B, f_C, Y). Through systematic uncertainty estimation analysis across 1,000 parameter combinations and varying statistical exposures, we demonstrate robust parameter estimation with proper uncertainty quantification in a highly correlated parameter space. Both methods provide parameter estimates with near-zero systematic bias and uncertainties limited only by statistics. Beyond the application to JUNO, this work establishes a methodological template for applying NLE across experimental neutrino physics for various problems where analytical likelihoods are intractable. As experiments like DUNE⁵⁵ and Hyper-Kamiokande⁵⁶ come online with increasingly stringent precision requirements, NLE (and SBI methods, in general) will become essential tools for extracting maximum physics information from data.

Results

Approach overview

As discussed above, the NLE approach requires models that can approximate the conditional probability density function (PDF) $p(x|\phi)$, where x represents the observations (the total number of photo-electrons $N_{p.e.}$ as a proxy for the visible energy) and ϕ denotes the parameters to be inferred. We present two complementary NLE implementations: TEDE and NFDE, both targeting the conditional density $p(x|\psi)$, where $\psi = \phi \cup S = \{k_B, f_C, Y, S\}$ combines the energy response parameters (ϕ) with the calibration source type (S). Incorporating S as an additional conditioning variable alongside ϕ enables us to utilize models in a unified way to handle different calibration sources during the parameter inference process, where all five calibration sources are fit simultaneously to jointly constrain the energy response parameters. Thus, the TEDE and NFDE models serve as fast likelihood estimators for the computationally expensive MC simulations.

While both approaches target the same conditional density estimation task, they differ in how they handle the problem. In the setting of the TEDE approach, we first approximate the conditional density using a histogram representation, then the TEDE model is trained to learn a direct mapping from ψ to the bin-to-bin approximations of the corresponding density. This approach naturally integrates with binned likelihood analyses. In contrast, NFDE directly models the continuous probability density through normalizing flows, enabling unbinned likelihood evaluation that preserves the full information content of the data. Below, we briefly outline the technical details of both models, while their comprehensive description can be found in the "Machine learning approach" subsection in Methods. Additionally, model training and hyperparameter optimization are described in the "Hyperparameter optimization and training" subsection in "Methods".

Specifically, the TEDE model uses a transformer encoder architecture to learn the mapping from input parameters ψ to histogram-based density approximations. The energy range, expressed in $N_{p.e.}$, from 400 to 16,400, is divided into $N_b = 800$ bins with a width of $\Delta x = 20 N_{p.e.}$. The model produces probability density values for each bin through a temperature-scaled softmax layer, which is given by:

$$\left\{ \hat{p}_i^{\text{TEDE}} \right\}_{i=1 \dots N_b} = \frac{1}{\Delta x} \left\{ \frac{e^{z_i/T}}{\sum_{j=1}^{N_b} e^{z_j/T}} \right\}_{i=1 \dots N_b}, \quad (1)$$

where z_i are the raw values from the preceding layer and T is the softmax temperature hyperparameter, which controls the sharpness of the output

distribution. By construction, this ensures proper normalization of the PDF: $\sum_{i=1}^{N_b} \hat{p}_i^{\text{TEDE}} \Delta x = 1$.

The NFDE model employs an exact approach to conditional density estimation based on normalizing flows. In the density estimation setting, the normalizing flows model is trained to transform a complex distribution into a well-known simple *base* distribution through a sequence of invertible transformations, *flows*. Usually, and in our specific case, the normal distribution is used as the base distribution. Invertibility of the flows allows transformation in both directions. The direction from the complex distribution to the base distribution is called *normalizing direction* (following the fact that the normal distribution is a common choice), while the opposite direction is called *generative direction*, since it allows us to produce new samples of the complex distribution by sampling from the base and passing the base samples through the model. The density estimation is based on the normalizing direction of the flow because it allows the exact likelihood computation for each value of the complex distribution. We train the NFDE model to transform the complex energy distributions of the calibration sources to the standard normal distribution through a sequence of invertible transformations f_θ , conditioned on ψ via dedicated neural network-based parameterizations. The exact probability density at any energy value x is then computed using the change of variables formula:

$$\hat{p}(x|\psi) = p_Z(f_\theta(x|\psi)) \left| \frac{\partial f_\theta(x|\psi)}{\partial x} \right|, \quad (2)$$

where p_Z is the density of the standard normal distribution and the Jacobian determinant, which is computed analytically, accounts for the volume change induced by the transformations.

The principle of this normalizing flow approach is illustrated in Fig. 1, which shows how the complex, multimodal energy distribution of $^{241}\text{Am-Be}$ is transformed through a sequence of 120 conditional flows into a simple Gaussian distribution. These transformations enable exact likelihood computation for unbinned data points—a key advantage for parameter estimation.

We evaluate the performance of both approaches in two stages. First, we demonstrate the ability of the models to accurately reproduce complex energy spectra across the parameter space. Then, we assess their effectiveness when used for parameter inference, evaluating the accuracy, precision, and robustness of the resulting energy response parameter estimations.

Our study is based on extensive MC data produced using the official JUNO software³⁷, containing approximately 1 billion simulated events across five calibration sources (^{137}Cs , ^{40}K , ^{60}Co , $^{241}\text{Am-Be}$, and $^{241}\text{Am-}^{13}\text{C}$) for different combinations of the parameters. Here, each event is ultimately represented by a single value of $N_{p.e.}$, which is conditioned on the parameters ψ . Following standard ML practices, we organize data into separate training (~600 million events), validation (~270 million events), and testing (~370 million events) datasets, with specific parameter grids designed to evaluate model performance in an unbiased manner. The training dataset covers a comprehensive grid of 9261 parameter combinations (21 uniformly spaced grid points per parameter) spanning $k_B \in [6, 24] \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$ ($\Delta k_B = 0.9 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$), $f_C \in [0, 1]$ ($\Delta f_C = 0.05$), and $Y \in [8000, 12000] \text{ MeV}^{-1}$ ($\Delta Y = 200 \text{ MeV}^{-1}$). Validation and testing datasets lie within the training parameter space, positioned exactly midway between the training points to evaluate models' interpolation performance in the worst-case scenario. We produce two types of validation datasets: validation dataset 1 covering a grid of 1,000 parameter points (with twice the training step size) and validation dataset 2 with high statistics at three specific points. Similarly, we construct two testing datasets: testing dataset 1 is a grid with 1000 different parameter combinations (with twice the training step size, shifted with respect to both validation and training grids), and testing dataset 2 is focused on a single parameter point (close to the center of the grid) with varying statistical exposures to perform a systematic uncertainty analysis. A detailed description of these datasets is provided in the “Data description” subsection in Methods.

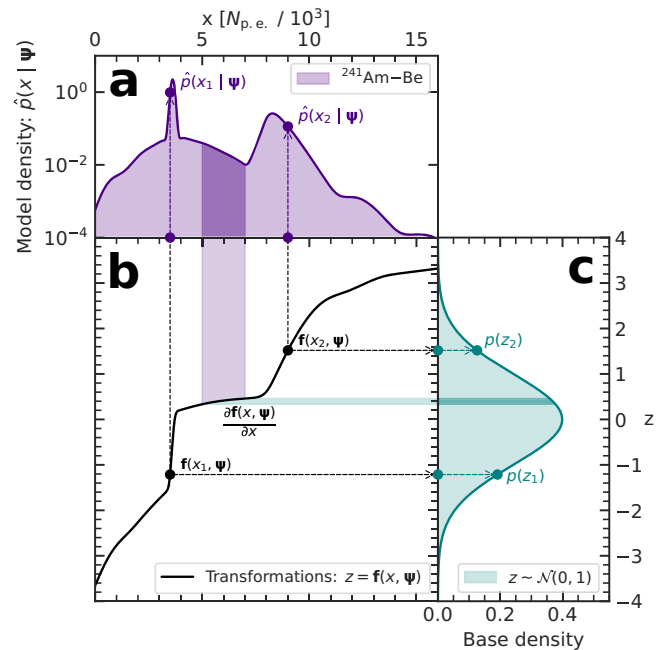


Fig. 1 | Illustration of the normalizing direction of the normalizing flows model. The Normalizing Flows Density Estimator (NFDE) transforms the complex energy distribution of $^{241}\text{Am-Be}$ into a simple Gaussian distribution through a series of learned invertible mappings $f_\theta : z = f_\theta(x, \psi) = f_K(f_{K-1}(\dots f_1(x, \psi)))$, where $\psi = \{k_B, f_C, Y, S\}$ is the vector of conditioning variables, and K is the number of flows. More specifically, NFDE transforms the values x (the visible energy expressed in number of photo-electrons, $N_{p.e.}$) (a) to latent space points z (c) through sequential flow transformations (b). The conditional probability density $\hat{p}(x|\psi)$ for any value of x given parameters ψ is calculated by applying the change-of-variables formula: $\hat{p}(x|\psi) = p_Z(z) \cdot |\partial f_\theta(x, \psi) / \partial x|$, where $p_Z(z)$ is the standard normal density and $\partial f_\theta(x, \psi) / \partial x$ is the Jacobian of the composite transformation f_θ . This approach enables exact likelihood evaluation needed for the unbinned parameter estimation procedure.

Models performance

To evaluate and compare the ability of both models to properly describe the data, we use a family of statistical distances based on the L_p -norm framework. Specifically, for probability distributions u and v with respective cumulative distribution functions U and V , we define the L_p -norm distance as follows:

$$d_p(u, v) = \left(\int_{-\infty}^{+\infty} |U(x) - V(x)|^p dx \right)^{1/p} \quad (3)$$

This formulation yields three classical statistical distances depending on the choice of p : Wasserstein ($p = 1$), Cramér-von Mises ($p = 2$), and Kolmogorov-Smirnov ($p = \infty$) distances. Using them as metrics, we quantify different aspects of the discrepancy between predicted and true probability distributions. Among them, we select the Cramér-von Mises distance (d_2) as our primary optimization criterion due to its balanced sensitivity to both global distribution shifts and local deformations, making it particularly useful for evaluating complex energy spectra with diverse features. Furthermore, this unified L_p -norm framework is suitable for both TEDE (with integral being approximated by a sum) and NFDE approaches. The “Metrics” subsection in Methods includes a comprehensive discussion of our metrics.

Figure 2 provides a visual comparison between the PDFs predicted by both models and the corresponding ground truth empirical PDF (limited in statistics) for all five calibration sources at a representative point from testing dataset 1 ($k_B = 10.05 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.225$, $Y = 11,700 \text{ MeV}^{-1}$). Both models demonstrate high fidelity in capturing diverse spectral features, including sharp peaks, Compton edges, and the continuous components characteristic of different source types. The TEDE model produces a binned

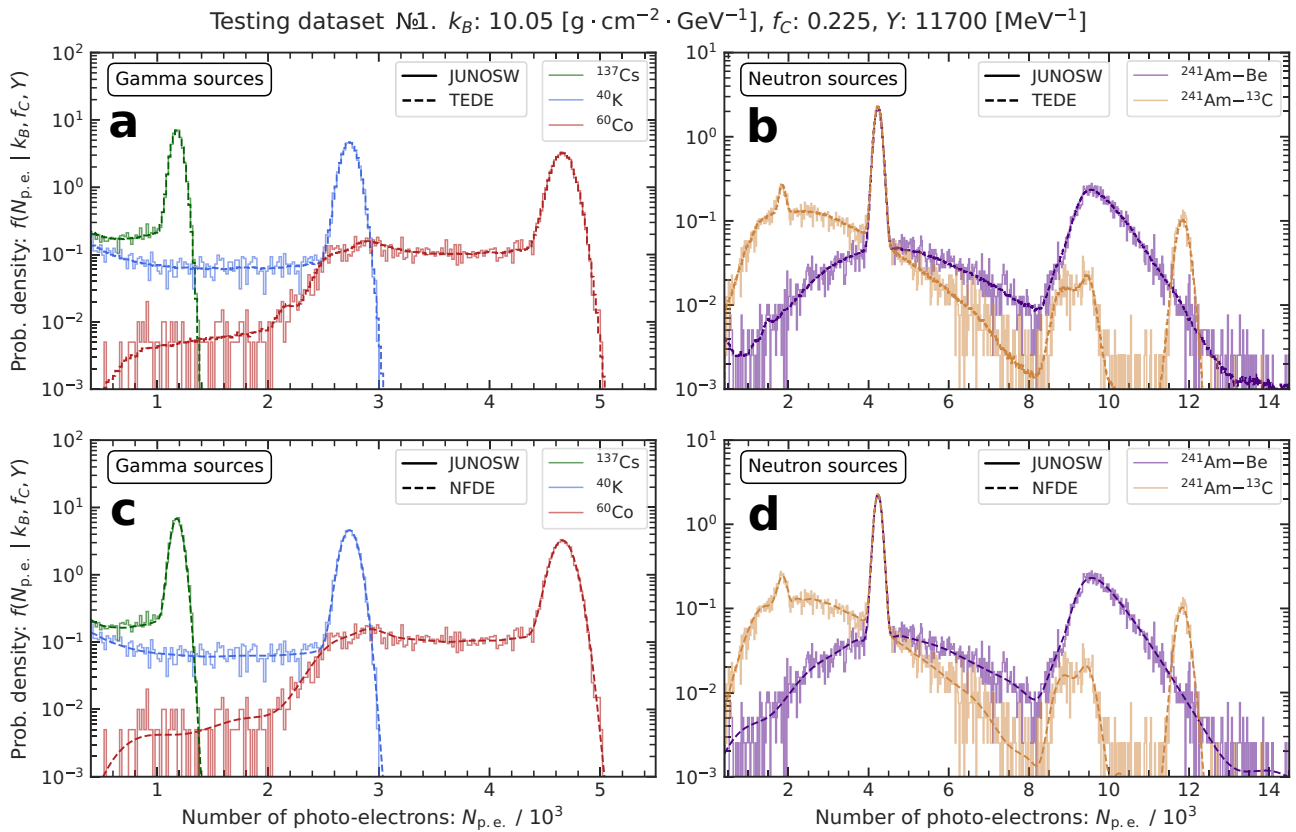


Fig. 2 | Comparison of the modeled probability density functions with the corresponding energy spectra at a representative point from testing dataset 1. All five calibration sources are shown with the following parameter values: $k_B = 10.05 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.225$, and $Y = 11,700 \text{ MeV}^{-1}$. Panels (a, c) show the gamma sources, while panels (b, d) show the neutron sources; this separation is for visualization purposes only. Panels (a–d) present the modeled energy spectra of each source built using the Transformer Encoder Density Estimator (TEDE) and the Normalizing Flows Density Estimator (NFDE), respectively. The histograms

represent empirical probability density functions from Monte Carlo simulations of the corresponding energy spectra, obtained using the JUNO software (JUNOSW) for the given parameter values. The statistical fluctuations in the number of events in each bin follow the expected Poisson distribution. Different calibration sources are represented by different colors: ^{137}Cs in green, ^{40}K in blue, ^{60}Co in red, $^{241}\text{Am-Be}$ in purple, and $^{241}\text{Am-}^{13}\text{C}$ in orange. Both models accurately reproduce the expected energy distributions of all sources, while NFDE presents smooth continuous curves.

Table 1 | Performance comparison between models

Statistical distance	Validation dataset		Testing dataset	
	TEDE	NFDE	TEDE	NFDE
$d_1 \times 10^{-3}$	5.0	5.7	4.5	5.4
$d_2 \times 10^{-3}$	2.6	2.6	2.4	2.6
$d_\infty \times 10^{-3}$	3.8	4.4	3.5	5.6

The table presents three statistical distances for both validation and testing datasets used as metrics in this study to compare the Transformer Encoder Density Estimator (TEDE) and Normalizing Flows Density Estimator (NFDE) models. Here, d_1 , d_2 , and d_∞ are the Wasserstein, Cramér-von Mises, and Kolmogorov-Smirnov distances, respectively. The metric values represent averages between the broad grid but sparse-statistics dataset (validation dataset 1) and three points high-statistics dataset (validation dataset 2), with similar averaging for the corresponding testing datasets. All values are scaled by 10^{-3} for readability. Lower values indicate better model performance in approximating the true probability distributions.

approximation that closely follows the target distribution while showing some insignificant residual fluctuations in the low-statistics regions, while the NFDE model generates smooth continuous curves that accurately represent the spectra PDFs, reproducing the very tiny spectrum features (such as the double gamma peak of the $^{241}\text{Am-}^{13}\text{C}$ at around $9000 N_{p.e.}$).

To quantitatively assess model performance, Table 1 presents a comparison of the statistical distance metrics for both models across validation and testing datasets. Both models achieve similar near-zero discrepancies with the data PDFs, validating both approaches as effective one-dimensional conditional density estimators. However, TEDE slightly outperforms NFDE

in all the metrics with differences within the same order of magnitude. While NFDE shows almost the same performance on the primary Cramér-von Mises (d_2) metric, it demonstrates noticeably worse results on Kolmogorov-Smirnov metric. This likely reflects the KS distance’s higher sensitivity to localized distribution differences and outliers, as it measures maximum deviation between cumulative distribution functions (CDFs) rather than integrated differences across the entire distribution. It is worth noting that when comparing these metrics, there is an additional source of uncertainty related to the methodological differences in CDF estimation between models: TEDE relies on binned histogram approximations, while NFDE employs unbinned analytical CDFs. This methodological difference may introduce small systematic variations in the distance measurements that are independent of the actual model performance.

Beyond accurately modeling individual spectra, the models must correctly capture the functional dependence of the spectrum shape on the energy response parameters. Figure 3 visualizes this capability for the NFDE model using the $^{241}\text{Am-}^{13}\text{C}$ source as an example. The heatmaps illustrate how the predicted conditional PDF evolves as each parameter is varied independently while holding the others fixed. The behavior of the modeled PDF varies smoothly and continuously across the parameter space, with no artificial discontinuities, demonstrating that the NFDE has learned physically consistent parameter dependencies essential for reliable interpolation during parameter estimation.

While the k_B and f_C parameters introduce a non-linear relationship between the deposited energy of an event and the collected charge (in $N_{p.e.}$), the dependence of the model’s output PDF values on the parameters for a




Fig. 3 | Energy response parameter impact on calibration energy spectra using the normalizing flows model. Heatmaps present the probability density given by the Normalizing Flows Density Estimator (NFDE) for the $^{241}\text{Am-}^{13}\text{C}$ neutron source as each parameter is varied independently: k_B (b), f_C (c), and Y (d). The rotated spectrum (a) shows the $^{241}\text{Am-}^{13}\text{C}$ spectrum modeled by NFDE for illustration. The following parameter values were used: $k_B = 6.0 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.5$, $Y = 10000 \text{ MeV}^{-1}$. For each panel, the two parameters that are not varied are fixed at

their median grid values. Color intensity (from black to yellow) represents conditional probability density function (PDF) values on a logarithmic scale, showing the expected parameter-specific effects. Note that the y-axis upper limit of $12.5k N_{p.e.}$ (number of photo-electrons) is chosen for visualization purposes only, as it encompasses the main parts of the spectra. A similar plot for the Transformer Encoder Density Estimator (TEDE) can be found in Supplementary Fig. 1.

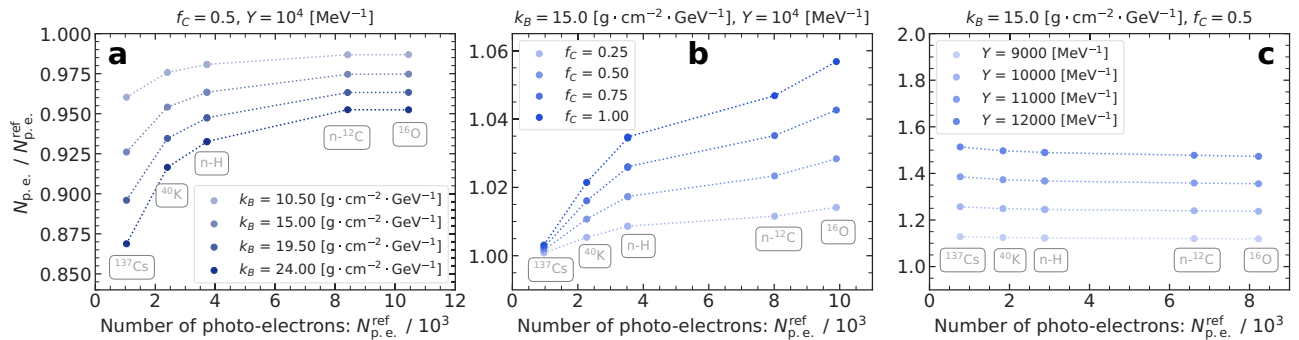


Fig. 4 | The effect of the energy response parameters as a function of energy using the normalizing flows model. This plot demonstrates how each parameter affects the relative positions of the main peaks across the energy scale of the calibration sources. The main peaks of the spectra from multiple calibration sources serve as standard candles: ^{137}Cs , ^{40}K , and neutron capture peaks (n-H, n- ^{12}C), as well as the ^{16}O excitation. Moving along the y-axis shows the same gamma peak under different energy response scenarios (different parameter values), while moving along the x-axis shows gamma peaks (i.e., different energies) under the same energy response

scenario. For each parameter (k_B, f_C and Y), five values spanning the full parameter range were selected while keeping the other two parameters fixed at the mean values of the grid. The y-axis shows the relative shift in peak positions compared to references. The plot also highlights the different influence of each parameter: k_B shows increasing non-linearity at lower energies (a), f_C exhibits greater effect at higher energies (b), while Y demonstrates a nearly constant scaling with slight deviations at low energies due to photomultiplier tube dark noise contributions (c). A similar plot for the transformer model can be found in Supplementary Fig. 2.

fixed deposited energy may or may not be linear. For example, in Fig. 3, the main gamma peaks shift as each parameter varies. For the k_B parameter, this shift is roughly linear, which is consistent with the first-order Taylor expansion of Birks' law being linear in k_B . For the f_C parameter, the behavior is more complex: there is almost no effect on the lower energy peaks, while the higher energy peaks show a more pronounced, approximately linear dependence on the parameter value. Additionally, the f_C parameter causes the peaks to become wider as the parameter increases, indicating a degradation in energy resolution. Finally, the dependence learned for the Y parameter is truly linear by the nature of the parameter. This describes how the PDF varies with the parameters at fixed deposited energy. It is important to distinguish this from the role those parameters play in producing the non-linear relationship between deposited energy and collected charge across the full energy scale when their values are fixed.

To analyze the non-linear relationship of energy response across different parameter values shown in Fig. 4, we use the following approach. First, we select the main peaks of the spectra as standard candles: ^{137}Cs , ^{40}K , n-H, n- ^{12}C peaks from neutron sources, as well as the ^{16}O peak of the $^{241}\text{Am-}^{13}\text{C}$ source. Note that we exclude the ^{60}Co photo-peak and peaks from

positron annihilation as they originate from the combined energy deposition of two simultaneous gamma emissions, thus effectively appearing more quenched. For each parameter, we take five values spanning the full range of the grid while keeping the other two parameters fixed at the mean values of the grid. We use spectra with the lowest parameter values as references. Using the NFDE model, we produce the PDFs for each parameter combination and measure the relative positions (in the number of photo-electrons) of each peak compared to the reference spectra. This analysis reveals different magnitudes of the parameter effects as shown in Fig. 4. Light yield (Y) demonstrates the strongest overall influence, followed by Birks' constant (k_B) and then by the Cherenkov light yield factor (f_C). The effect of k_B is distinctly non-linear, becoming more pronounced at lower energies. Conversely, f_C shows a non-linear behavior with its effect magnitude increasing at higher energies while having minimal impact in the low-energy region and almost zero impact for the ^{137}Cs source. This makes the ^{137}Cs source particularly valuable for constraining the other parameters. The light yield response appears nearly constant, with slight deviations at low energies attributable to PMT dark noise contributions, whose relative impact decreases as energy increases.

Though relatively subtle, these non-linear effects are critical for the high precision determination of the detector energy scale (as in the case of JUNO³⁶). Using the TEDE model, we obtained similar results to the ones presented for NFDE in Figs. 3 and 4. These corresponding TEDE results are provided in Supplementary Figs. 1 and 2.

The models' ability to learn these complex dependencies (including their non-linear nature) directly from the data is essential for their effectiveness in the subsequent parameter estimation task, particularly for breaking degeneracies between highly correlated parameters. As evidenced by the parameter recovery accuracy shown in the subsequent section, both models successfully capture these small but systematic non-linear effects.

The effectiveness of the NFDE approach can be further assessed by examining how well it performs its core function: transforming complex data distributions into the standard normal distribution. Supplementary Fig. 3 shows an example of the model's success in this task for a point from validation dataset 2. The visualization shows how the complex, multimodal energy distributions of the five calibration sources (a) are progressively transformed through the learned flows (b) into well-formed standard normal distributions (c). This successful normalization across different source types is persistent across the entire parameter space of the validation datasets and thus confirms that the model has learned appropriate conditional transformations for all calibration sources within the parameter space.

Parameter estimation

Single representative point in the parameter space. To demonstrate the parameter recovery capabilities of the models, we first examine the representative example from testing dataset 1, corresponding to true parameters $k_B = 10.05 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.225$, $Y = 11,700 \text{ MeV}^{-1}$. For each point and each source of testing dataset 1, we generate 10,000 events. Due to the decay scheme of the neutron sources, where both a neutron and de-excitation gammas are emitted, the final amount of events is doubled for these sources. Additionally, events with total charge collected below the trigger threshold (of around 200 photo-electrons) are discarded in the simulation. This effect varies by source: from several percent event loss for ¹³⁷Cs to almost no event loss for ⁶⁰Co, with the total event loss for all five sources combined being about 2.1%. Thus, the total statistics for each point of testing dataset 1 are variable but approximately equal to 70,000 events. For simplicity, we define the term "statistical exposure level" as the number of generated events per source (i.e., 10,000 for testing dataset 1), which is therefore approximately seven times smaller than the actual total statistics of the combined sample.

We perform a joint fit for all five calibration sources using likelihood functions specific to each model's approach. The total log-likelihood function $\ln\mathcal{L}(\mathcal{D} | \phi)$ is defined as a sum of individual source log-likelihoods, $\ln\mathcal{L}_s(\mathcal{D}_s | \phi)$:

$$\ln\mathcal{L}(\mathcal{D} | \phi) = \sum_{s=1}^5 \ln\mathcal{L}_s(\mathcal{D}_s | \phi), \quad (4)$$

where s is the source index, \mathcal{D} is the total dataset, \mathcal{D}_s is the dataset for a specific source S , and $\phi = \{k_B, f_C, Y\}$ are the three energy response parameters. In addition to ϕ , the fit simultaneously determines normalization constants $N_{s=1, \dots, 5}$ for each of the five calibration sources: ¹³⁷Cs, ⁴⁰K, ⁶⁰Co, ²⁴¹Am-Be, and ²⁴¹Am-¹³C, respectively. These N_s parameters represent the expected total number of events and are treated as free parameters.

For TEDE, we use a binned Poisson log-likelihood, which naturally accommodates its histogram-based representation:

$$\begin{aligned} \ln\mathcal{L}_s(\mathcal{D}_s | \phi) &= \\ &= \sum_{i=1}^{N_b} \left[m_{s,i} \cdot \ln \mu_{s,i}(\phi) - \mu_{s,i}(\phi) - \ln(m_{s,i}!) \right], \end{aligned} \quad (5)$$

where $m_{s,i}$ is the observed number of events for each of the N_b bins and $\mu_{s,i}(\phi) = N_s \cdot \hat{p}_i^{\text{TEDE}}(\phi, S)$ is the expected number of events, with

$\hat{p}_i^{\text{TEDE}}(\phi, S)$ being the probability density per bin predicted by the TEDE model. During optimization, the constant term $\ln(m_{s,i}!)$ is omitted since it does not depend on ϕ .

For NFDE, we employ an extended unbinned likelihood^{57,58}, which benefits from the model's ability to evaluate exact probability densities at individual data points $\{x_{s,i}\}_{i=1}^{M_s}$:

$$\begin{aligned} \ln\mathcal{L}_s(\mathcal{D}_s | \phi) &= \\ &= M_s \ln N_s - N_s - \ln(M_s!) + \sum_{i=1}^{M_s} \ln \hat{p}(x_{s,i} | \phi, S), \end{aligned} \quad (6)$$

where M_s is the total number of data points and $\hat{p}(x_{s,i} | \phi, S)$ are the point-by-point density estimates provided by the NFDE model. Analogously to $\ln(m_{s,i}!)$ in the binned case, the $\ln(M_s!)$ term is omitted during optimization since it does not depend on ϕ .

The Bayesian nested sampling algorithm⁵¹ (from `ultranest`⁵⁰), which efficiently explores the parameter space and computes posterior distributions even for complex, multimodal cases, was employed to infer the energy response parameters for both models.

The posterior distributions shown in Fig. 5 reveal well-constrained parameters, with estimates that are both precise and accurate. Using multiple calibration sources helps effectively reduce the uncertainties in energy response modeling. Both models achieve comparable precision, successfully recovering the input parameters within their uncertainties, as summarized in Table 2.

As described above, the testing points lie midway between training grid points (visible as gray lines in Fig. 5), i.e., $0.45 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$ away in k_B , 0.025 away in f_C , and 100 MeV^{-1} away in Y from the nearest training points. This demonstrates the models' ability to interpolate effectively in this worst-case scenario—a critical feature for practical applications where the true parameters may lie anywhere within the parameter space.

The posterior distribution plots also reveal the correlations between parameters, which are quantitatively presented in Supplementary Fig. 4 for the NFDE fit. Such high-correlation behavior is expected given parameters' joint influence on the non-linear effects of the energy response. While the correlation structure shown in Supplementary Fig. 4 is representative, the analysis of all fits across both testing datasets shows consistent correlation patterns throughout the parameter space. The fact that unbiased parameter estimates were achieved with proper uncertainty coverage (as demonstrated in the subsequent subsection) further validates that the models correctly capture these correlations. The Bayesian nested sampling approach provides a complete picture of the parameter space, including these correlations, degeneracies, and constraints.

Aggregated results and the robustness analysis. To assess the robustness and potential biases of our approach, we conducted a comprehensive analysis across varying statistical conditions and throughout the parameter space.

Stability under varying statistical exposure levels. The analysis based on testing data 2 corresponds to results obtained for a single point in the (k_B, f_C, Y) parameter space with varying statistical exposure levels. This point is set at $k_B = 15.45 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.525$, and $Y = 10,100 \text{ MeV}^{-1}$, positioned near the center of the parameter space and midway between training grid points.

Figure 6 presents a comparative assessment of the models, TEDE and NFDE, studying their performance at a single point in the parameter space across varying statistical exposure levels. Each data point corresponds to 1000 independent fits performed on different Poisson realizations of the expected distribution (each generated using a different random seed in the simulator), with a total of 10,000 fits conducted across five exposures and two models. One key observation is that both models exhibit a near-zero median bias (colored solid lines), with relative deviations below 0.25% for k_B , 0.55% for f_C , and 0.02% for Y , calculated with respect to their true values. In addition, the medians of the estimated uncertainties (black solid lines)

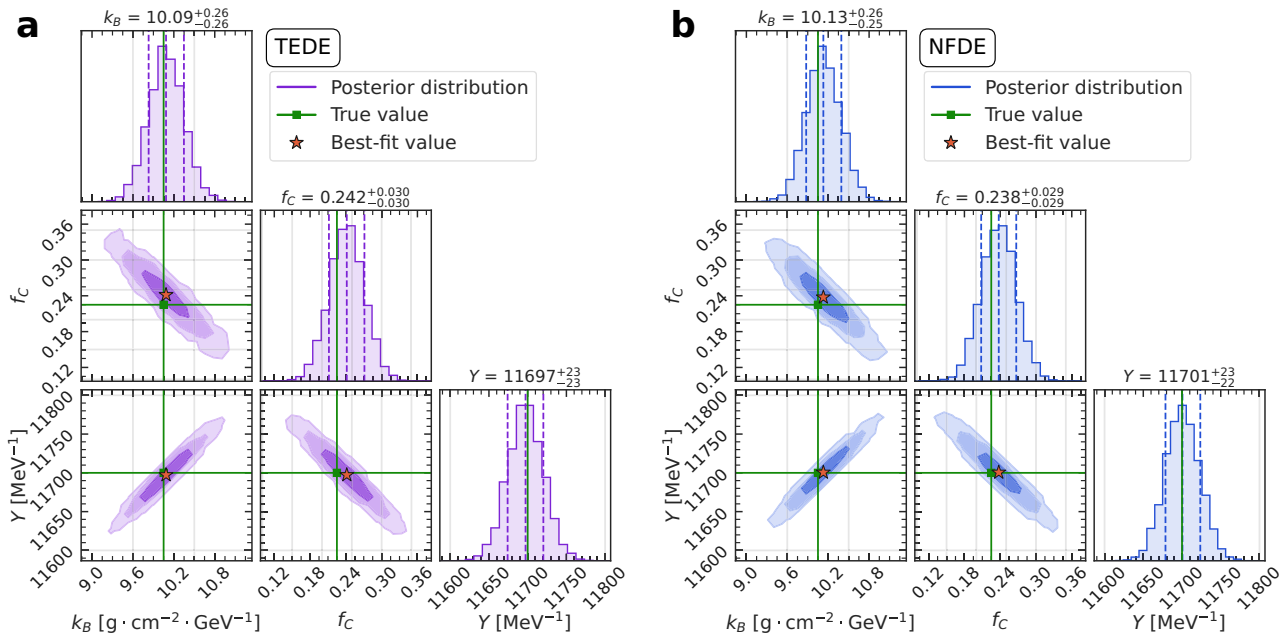


Fig. 5 | Parameter estimation performance for a single representative point of testing dataset 1. Corner plots show the one- and two-dimensional marginalized posterior probability distributions for the energy response parameters (k_B, f_C, Y), obtained using the Transformer Encoder Density Estimator (TEDE, **a**) and Normalizing Flows Density Estimator (NFDE, **b**) models with the Bayesian nested

sampling algorithm (from *ultranest*). Green lines indicate the true parameter values, while orange stars show the best-fit values. The gray grid represents the parameter values of the training dataset. Both models yield posterior distributions centered near the true values, demonstrating the effective parameter recovery through the combined fit across five calibration sources.

Table 2 | Parameter estimation results for a single representative point from testing dataset 1

Parameter	True value	Estimated value	
		TEDE	NFDE
$k_B, \text{g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$	10.05	$10.09^{+0.26}_{-0.26}$	$10.13^{+0.26}_{-0.25}$
f_C	0.225	$0.242^{+0.030}_{-0.030}$	$0.238^{+0.029}_{-0.029}$
Y, MeV^{-1}	11,700	$11,697^{+23}_{-23}$	$11,701^{+23}_{-22}$

The table presents a comparison of true values of the energy response parameters and best-fit estimates (with their corresponding uncertainties, computed as the 34.1% of the probability on either side of the median). Parameter inference is performed using the Bayesian nested sampling method from *ultranest* for both the Transformer Encoder Density Estimator (TEDE) and Normalizing Flows Density Estimator (NFDE) models.

consistently align with the actual 1σ widths of the distributions of best-fit values (shaded regions). Here, the 1σ widths represent intervals that enclose 34.1% of the probability on either side of the median in the best-fit distributions. The reported medians of the estimated uncertainties, along with the medians of the lower and upper 1σ bounds, are derived from 1,000 individual fits. These findings indicate that neither model systematically skews parameter estimates, reinforcing confidence in the overall accuracy of the best-fit estimates. Additionally, the fit results for TEDE and NFDE (black lines on Fig. 6) show negligible differences in the uncertainty estimates, highlighting alignment of the results not only within each model but also between them.

Moreover, for both models, the estimated uncertainties closely follow the expected dependence on the number of events per source, as dictated by pure statistical errors³⁷. Specifically, the uncertainties scale as $1/\sqrt{N}$, where N is the statistical exposure level. Thus, the proportionality constants are:

$$\frac{0.32 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}}{\sqrt{N/10^4}}, \frac{0.026}{\sqrt{N/10^4}}, \frac{21 \text{ MeV}^{-1}}{\sqrt{N/10^4}},$$

for k_B, f_C , and Y , respectively, with the constants extracted by fitting the squared median uncertainties as a linear function of $1/N$. At the benchmark

exposure level of 10^4 , this translates to relative uncertainties of approximately 2.1% for k_B , 4.9% for f_C , and 0.2% for Y .

Finally, the 1σ widths of the estimated uncertainty distributions (error bars of the corresponding points of black solid lines) are rather small for both models, indicating fit stability and reliability of obtained best-fit values and their estimated errors.

JUNO’s calibration program is expected to collect significantly higher statistics, with employed sources typically reaching activities of ~ 100 Hz and detailed calibrations lasting hours⁴⁵. Thus, the experimental data will provide sufficient events to reach or exceed the parameter precision demonstrated in this study. Once best-fit values are determined, model-related systematic uncertainties can be re-evaluated by applying the strategy outlined in this section with a focused grid around the best-fit point.

These findings demonstrate that both TEDE and NFDE models are robust, exhibiting negligible systematic bias and producing uncertainty estimates that align well with expected statistical uncertainty. Their accuracy remains consistent across different exposure levels, and uncertainties scale precisely with the number of events per source, confirming their reliability for statistical analyses.

Stability across varying parameter space. This analysis is based on testing dataset 1 and investigates the behavior of the TEDE and NFDE models under a fixed statistical exposure (10,000 events per source) while systematically varying the parameter space. Each data point in Fig. 7 represents the aggregated results of 100 individual fits, allowing for a detailed examination of bias behavior across the parameter range. The number 100 reflects our parameter grid structure: when analyzing the bias for a specific parameter value (e.g., one value of k_B), we aggregate results from 100 fits corresponding to all combinations of the other two parameters (10 values of $f_C \times 10$ values of Y) at that fixed k_B .

Both TEDE and NFDE models demonstrate biases that consistently lie within the 1σ best-fit width, indicating that neither model exhibits significant systematic deviations. This suggests that the models are capable of successfully interpolating between points in the training grid, maintaining accuracy and robustness across the parameter space. Furthermore, the observed bias behavior reinforces confidence in their predictive reliability, ensuring stable performance even under varying parameter conditions.

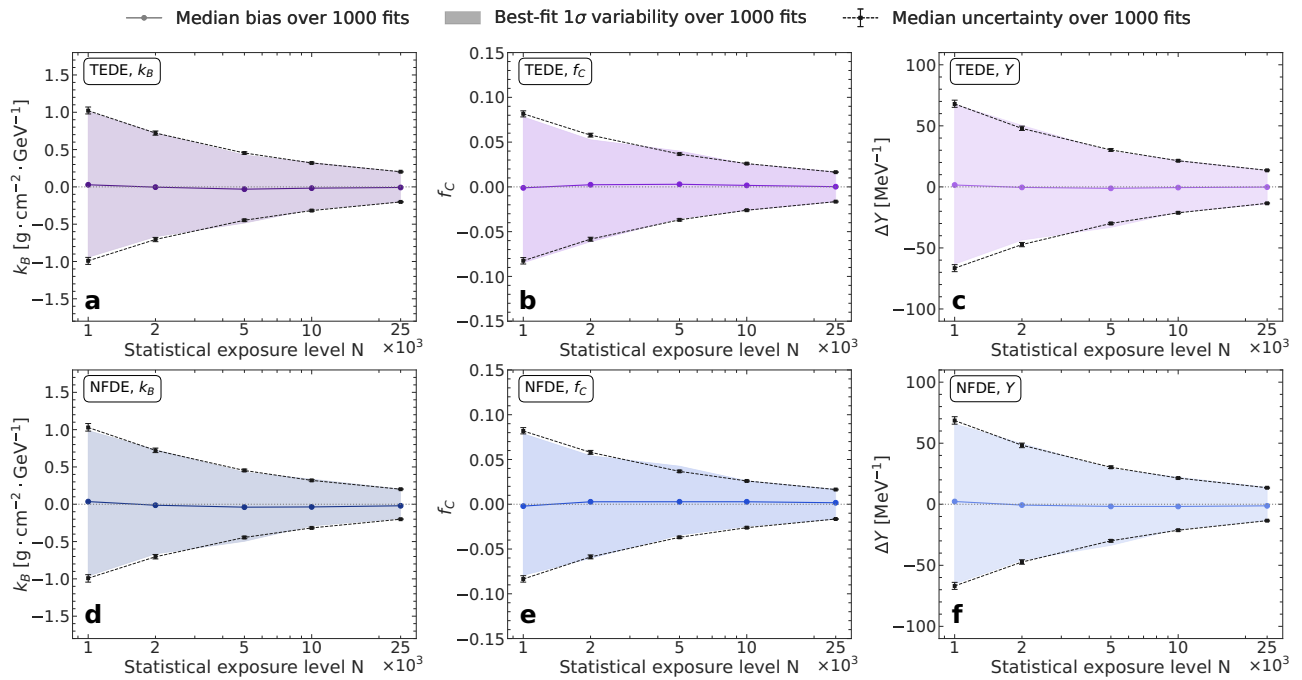


Fig. 6 | Aggregated parameter estimation results for testing dataset 2. The plots show aggregated results for k_B (a, d), f_C (b, e), and Y (c, f), including the median bias, the 1σ widths of the best-fit value distributions, and the estimated uncertainties. Results are shown for both the Transformer Encoder Density Estimator (TEDE, a–c) and Normalizing Flows Density Estimator (NFDE, d–f). Solid lines with circle markers represent the median bias (i.e., the median of the distribution of differences between best-fit and true parameter values). Purple and blue shaded regions indicate the 1σ widths of the best-fit value distributions (i.e., 68.3% coverage of the obtained

best-fit values). Dashed lines with square markers represent the medians of the 1σ uncertainty widths estimated by the fits. Finally, the black error bars show the 1σ widths of these estimated uncertainty distributions. Each point corresponds to results from 1000 fits at a given exposure level. These results demonstrate the robustness and accuracy of both models across varying statistical scenarios, showing near-zero biases and estimated uncertainties that closely align with the actual variability in best-fit values.

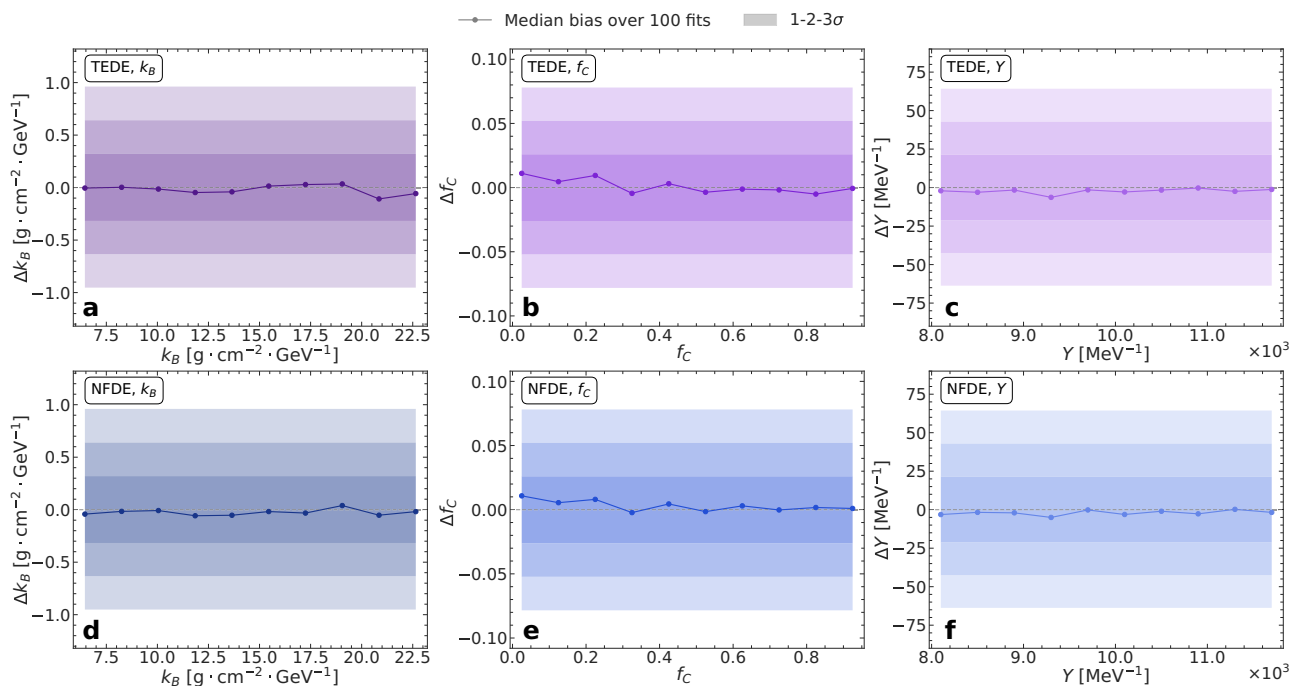


Fig. 7 | Performance across the considered parameter space, obtained for testing dataset 1. Results are shown for both the Transformer Encoder Density Estimator (TEDE, a–c) and Normalizing Flows Density Estimator (NFDE, d–f). The dataset comprises 1000 points sampling the (k_B, f_C, Y) space, each with the exposure level of 10,000. Purple and blue shaded areas represent 1, 2, and 3σ uncertainties obtained

from testing dataset 2 for the same exposure level. Each point represents the median bias, defined as the difference between the best-fit and true parameter values, computed over 100 fits. Both models maintain near-zero bias across the parameter space, demonstrating their robustness and stability.

Differences between the parameter minimization methods. In addition to the Bayesian nested sampling algorithm, we employed two other methods for parameter estimation: the Frequentist MIGRAD optimization algorithm with its MINOS uncertainty estimation (from `iminuit`⁵⁹), and a self-written Bayesian Metropolis-Hastings Markov chain Monte Carlo (MCMC) method^{60,61}. All three methods are described in detail in the “Parameter estimation” subsection in Methods.

The analysis demonstrated the reliability of both TEDE and NFDE models across different parameter estimation methods. Frequentist approaches (MIGRAD+MINOS) and Bayesian techniques (nested sampling and Metropolis-Hastings MCMC) produce consistent best-fit values and confidence or credible intervals, ensuring robustness in parameter inference. While MIGRAD+MINOS excels in speed, making it ideal for rapid estimations, Bayesian methods provide a more comprehensive statistical understanding by delivering full posterior distributions and offering deeper insights into parameter behavior. The results obtained with MIGRAD+MINOS and Metropolis-Hastings MCMC can be found in Supplementary Figs. 5–8, respectively.

This balance between efficiency and depth highlights the strengths of each approach, underscoring their applicability in diverse analysis scenarios.

Discussion

NFDE’s ability to produce continuous PDFs for unbinned analysis provides a potential advantage in statistical precision for parameter recovery, particularly in low-statistics cases. Although in our comparative analyses this difference is negligible, it could be more significant in other applications. Nevertheless, TEDE offers computational efficiency and implementation simplicity while maintaining the same performance, serving as a complementary approach that validates the overall methodology. We would also like to highlight that computational considerations do not limit the adoption of unbinned likelihood methods nowadays. While traditionally viewed as computationally expensive, our implementation of the unbinned analyses can be efficiently parallelized on both GPU and CPU with up to M parallel units (where M is the number of events). The increased availability of modern computing resources addresses a long-standing barrier to wider adoption of unbinned techniques in large-scale experiments. The scalability of these approaches is also promising when considering extensions to higher-dimensional parameter spaces or larger data volumes. Our current implementation of NFDE scales linearly with data size, while the parameter space complexity primarily affects the training phase rather than inference.

To provide quantitative context for the computational advantages of our approach, we present timing comparisons for the key components of parameter estimation. Our TEDE model takes approximately 3 ms to build each PDF, independently of the number of events in the dataset. The NFDE model requires approximately $40 + 0.002 \cdot M$ ms for probability density estimations, corresponding to 60 ms for a typical dataset of 10^4 events. In contrast, the full event simulation and reconstruction, using the official JUNO software, takes approximately 15 hours per parameter point per CPU for 10^4 events. Critically, this number of events is insufficient to construct reliable PDFs with negligible statistical fluctuations, necessitating orders of magnitude more events and correspondingly longer computation times. This represents a computational speedup of several orders of magnitude (in $\sim 10^6$ times in the case of 10^4 events), enabling previously intractable analyses such as comprehensive uncertainty quantification.

Despite using neural networks, both approaches maintain physical interpretability—an important advantage over typical “black box” ML models⁶². The explicit connection to the underlying physical parameters (k_B , f_C , Y) ensures that the results remain physically meaningful and that uncertainties directly translate to constraints on detector properties.

This work also represents one of the relatively rare successful applications of one-dimensional conditional normalizing flows in an experimental physics context. While normalizing flows have gained significant attention in particle physics for tasks such as event generation, anomaly detection, and variational inference in high-dimensional settings⁶³, their use for one-dimensional conditional density estimation in particle physics has been limited.

We also demonstrate that the presented approach is capable of resolving highly correlated parameters. The energy response parameters in our study exhibit strong correlations (as shown in Supplementary Fig. 4), presenting a challenging parameter estimation scenario. Both models successfully capture this correlation structure.

A fundamental assumption of our method is that the MC simulation accurately represents the underlying physics and the detector. As common in particle physics, JUNO’s simulation is based on well-validated `Geant4` physics models and detailed detector descriptions, but any systematic discrepancies between simulation and experimental data (mismodeled physics processes, imperfect geometry description, etc.) may lead to biased parameter estimates. Some of these effects can be incorporated a posteriori if they can be modeled analytically at the level of the parameter estimation procedure, while others may require prior modifications to the simulation framework and retraining the models from scratch on updated data before parameter estimation. This limitation is inherent to any MC tuning approach and highlights that MC tuning is a refinement step that requires a fundamentally sound simulation as its foundation, not a remedy for underlying modeling deficiencies.

Looking forward, several avenues for improvement exist. Higher-dimensional normalizing flows could potentially model the joint distribution of energy and other observables (such as time distribution or particle identification discriminators). Although directly extending to higher dimensions can be challenging, methods such as those presented in ref. 64 have the potential to effectively address this issue. Prior constraints could also be incorporated into the Bayesian framework to further enhance parameter estimation when prior knowledge is available and justified to be used⁶⁵. Extensions to time-dependent calibration, where detector parameters evolve, represent another promising direction for future research.

Methods

Data description

The well-established approach in the field of machine learning is to prepare three different categories of datasets: so-called training, validation, and testing datasets. The training dataset is used to train a model (to adjust the model’s learnable weights by minimizing a loss function). The validation dataset is used to evaluate model performance *during* the training process and guides hyperparameter optimization, while the testing dataset provides a final unbiased quantification of model performance once it is *trained*. These datasets must remain independent from one another to prevent biases and overfitting.

The data used in this study were generated using the official JUNO software³⁷, based on a detailed `Geant4`^{66,67} detector simulation. Different processes during photo-electron collection and further current amplification (e.g., dark noise), as well as intrinsic effects of electronics, are also simulated. The waveform reconstruction is performed using the deconvolution method implemented in the official JUNO software based on the method described in ref. 68. Our training, validation, and testing datasets are constructed to balance comprehensive coverage of the parameter space with the density of the points. We simulate three gamma sources (¹³⁷Cs, ⁴⁰K, ⁶⁰Co) and two neutron sources (²⁴¹Am-Be, ²⁴¹Am-¹³C) placed at the detector center.

Training dataset. We define a (k_B, f_C, Y) grid with 21 uniformly spaced values for each parameter. The ranges are chosen to encompass the default values currently used in the JUNO software ($k_B = 12.05 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}$, $f_C = 0.517$, $Y = 9846 \text{ MeV}^{-1}$)³⁶ with sufficient margin to robustly account for expected parameter variations. The grid is defined as follows:

$$\Theta_{k_B}^{\text{Train}} = \{6 + 0.9j\} \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1},$$

$$\Theta_{f_C}^{\text{Train}} = \{0 + 0.05j\},$$

$$\Theta_Y^{\text{Train}} = \{8000 + 200j\} \text{ MeV}^{-1},$$

$$\text{where } j : j \in \mathbb{Z}, 0 \leq j \leq 20.$$

The full training parameter space is then defined as $\Theta^{\text{Train}} = \Theta_{k_B}^{\text{Train}} \times \Theta_{f_C}^{\text{Train}} \times \Theta_Y^{\text{Train}}$, resulting in $21^3 = 9261$ unique parameter combinations. For each parameter point and each of the five calibration sources, we simulate 10^4 events. Events with light collection below the trigger threshold (~ 200 photo-electrons) are discarded. For the neutron sources ($^{241}\text{Am-Be}$, $^{241}\text{Am-}^{13}\text{C}$), simulations yield prompt (de-excitation gamma and neutron recoil) and delayed (gamma from neutron capture) events, both of which are used together in a single distribution. Therefore, to construct datasets for model training, ~ 600 million events are simulated.

Validation datasets. We produce two validation datasets for monitoring model performance during training and optimizing hyperparameters:

1) *Validation dataset 1.* This dataset uses a $10^3 = 1000$ point grid, shifted relative to the training grid to lie exactly midway between training points, probing the models' interpolation capability in an unbiased manner under the most difficult conditions:

$$\begin{aligned} \Theta_{k_B}^{\text{Val}_1} &= \{7.35 + 1.8j\} \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, \\ \Theta_{f_C}^{\text{Val}_1} &= \{0.075 + 0.1j\}, \\ \Theta_Y^{\text{Val}_1} &= \{8300 + 400j\} \text{ MeV}^{-1}, \\ &\text{where } j : j \in \mathbb{Z}, 0 \leq j \leq 9. \end{aligned}$$

Thus, $\Theta^{\text{Val}_1} = \Theta_{k_B}^{\text{Val}_1} \times \Theta_{f_C}^{\text{Val}_1} \times \Theta_Y^{\text{Val}_1}$. Analogously to the training dataset, we simulate 10^4 events per source per point, yielding ~ 70 million events.

2) *Validation dataset 2.* This dataset comprises three specific parameter points chosen from the validation dataset 1 grid space, but simulated with high statistics (10^7 events per source per point, ~ 200 million total events):

$$\begin{aligned} \Theta^{\text{Val}_2} &= \{(k_{B,i}, f_{C,i}, Y_i) | i = 1, 2, 3\} \\ &= \{(7.35 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, 0.075, 8300 \text{ MeV}^{-1}), \\ &\quad (14.55 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, 0.475, 9900 \text{ MeV}^{-1}), \\ &\quad (23.55 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, 0.975, 11900 \text{ MeV}^{-1})\} \end{aligned}$$

Although this dataset covers only three different points, the significantly higher statistics per spectrum (1000 times greater than validation dataset 1) provide statistically precise ground truth spectra, enabling accurate evaluation of model performance with minimal statistical uncertainties.

Testing datasets. Two testing datasets, independent from the training and validation ones, are used for final model evaluation.

1) *Testing dataset 1.* Similar to validation dataset 1, this uses a $10^3 = 1000$ point grid, again shifted to lie midway between training points, but different from the validation grid:

$$\begin{aligned} \Theta_{k_B}^{\text{Test}_1} &= \{6.45 + 1.8j\} \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, \\ \Theta_{f_C}^{\text{Test}_1} &= \{0.025 + 0.1j\}, \\ \Theta_Y^{\text{Test}_1} &= \{8100 + 400j\} \text{ MeV}^{-1}, \\ &\text{where } j : j \in \mathbb{Z}, 0 \leq j \leq 9. \end{aligned}$$

Thus, $\Theta^{\text{Test}_1} = \Theta_{k_B}^{\text{Test}_1} \times \Theta_{f_C}^{\text{Test}_1} \times \Theta_Y^{\text{Test}_1}$. As with the other grid-based datasets, we simulate 10^4 events per source per point (~ 70 million total). This dataset is used to assess the uniformity of parameter estimation performance across the entire parameter space, see Fig. 7.

2) *Testing dataset 2.* This dataset focuses on a single parameter point near the center of the parameter space ($\{15.45 \text{ g} \cdot \text{cm}^{-2} \cdot \text{GeV}^{-1}, 0.525, 10,100 \text{ MeV}^{-1}\}$), also located midway between training points. For this point, we produce 1000 independent simulations (using different random seeds) for each of five different event exposures per source: $\{10^3, 2 \times 10^3, 5 \times 10^3, 10^4, 2.5 \times 10^4\}$ events. This extensive dataset (~ 300 million total events) allows us to study how parameter uncertainties and biases

evolve as a function of event statistics. Moreover, it enables a comparison between the uncertainty estimates derived from individual fits and the actual distribution of best-fit values, therefore validating the uncertainty estimation procedure and evaluating potential systematic effects, see Fig. 6.

The spatial relationship between these different parameter grids is shown in Supplementary Fig. 9, which presents a two-dimensional projection of the k_B - Y parameter space. This structured approach to dataset design allows us to comprehensively evaluate model performance: training on a dense grid while testing on points that lie exactly midway between training points creates the most challenging conditions for interpolation, ensuring that our evaluation measures the models' ability to generalize to unseen regions of the parameter space rather than simply memorize training examples. The carefully arranged validation and testing grids, offset from each other and from the training grid, enable robust cross-validation and ensure unbiased performance assessment. Finally, we note that our study is based on uniformly sampled parameter grids, which represent a natural and effective sampling strategy for MC tuning applications where training data generation is under full control. A uniform grid minimizes the average distance to the nearest training point, ensuring reliable coverage of the parameter space.

Machine learning approach

Directly using the full chain of the JUNO software within the inference frameworks (see the "Parameter estimation" subsection below) for energy response parameter tuning is computationally prohibitive due to the extensive time required for producing sufficient event statistics across the parameter space. To overcome this bottleneck, we develop neural likelihood estimators designed to emulate the simulation chain. These models provide rapid estimates of the conditional PDF of the energy spectrum, $p(x|\psi)$, where x represents the visible energy (in number of reconstructed photo-electrons $N_{p.e.}$) and $\psi = \{k_B, f_C, Y, S\}$ is the vector of conditioning variables comprising the energy response parameters (k_B, f_C, Y) and the calibration source type (S). Once trained using the simulation dataset, these fast surrogate models replace the full simulation during the parameter estimation procedures.

Thus, the core task for the ML models is formulated as one-dimensional conditional density estimation. This task becomes challenging due to the complex and multi-modal nature of the data. We investigate two different approaches as powerful tools to achieve the goals: (1) a direct multi-output regressor and (2) a conditional normalizing flows model. While both of our approaches aim to approximate the conditional densities, the regressor approach provides an estimation through representing the density function using a histogram, and the normalizing flows solution offers the exact likelihood modeling. These strategies require different model architectures, training details, and likelihood formulations for subsequent parameter estimation, as detailed below.

Transformer Encoder Density Estimator (TEDE)

The method. This direct regression approach begins by approximating the one-dimensional target continuous PDF $p(x|\psi)$ with a histogram using a fixed binning scheme. Let this histogram-based approximation, derived from the JUNO software output for a given condition ψ , be represented by the vector $\mathbf{p}^{\text{JSW}} = \{p_1^{\text{JSW}}, p_2^{\text{JSW}}, \dots, p_{N_b}^{\text{JSW}}\}$. Here, JSW stands for "JUNO software". Each element p_i^{JSW} corresponds to the probability density in the i -th bin, calculated as $p_i^{\text{JSW}} = n_i / (N_{\text{tot}} \cdot \Delta x)$, where n_i is the event count in bin i , N_{tot} is the total number of events for that condition, and Δx is the bin width. By definition, this target distribution is normalized such that its discrete integral equals one: $\sum_{i=1}^{N_b} p_i^{\text{JSW}} \cdot \Delta x = 1$. The TEDE model, denoted as h_θ with learnable parameters θ , is then trained to learn the direct mapping from the input conditions ψ to the vector of estimated PDF values $\hat{\mathbf{p}}^{\text{TEDE}} = \{\hat{p}_1^{\text{TEDE}}, \hat{p}_2^{\text{TEDE}}, \dots, \hat{p}_{N_b}^{\text{TEDE}}\}$ for the corresponding energy bins:

$$\hat{\mathbf{p}}^{\text{TEDE}} = h_\theta(\psi), \tag{7}$$

where $\sum_{i=1}^{N_b} \hat{p}_i^{\text{TEDE}} \cdot \Delta x = 1$ must hold.

We discretize the x space into $N_b = 800$ bins with a bin width of $\Delta x = 20 N_{p.e.}$, spanning from 400 to 16400 $N_{p.e.}$. The core of the model is a

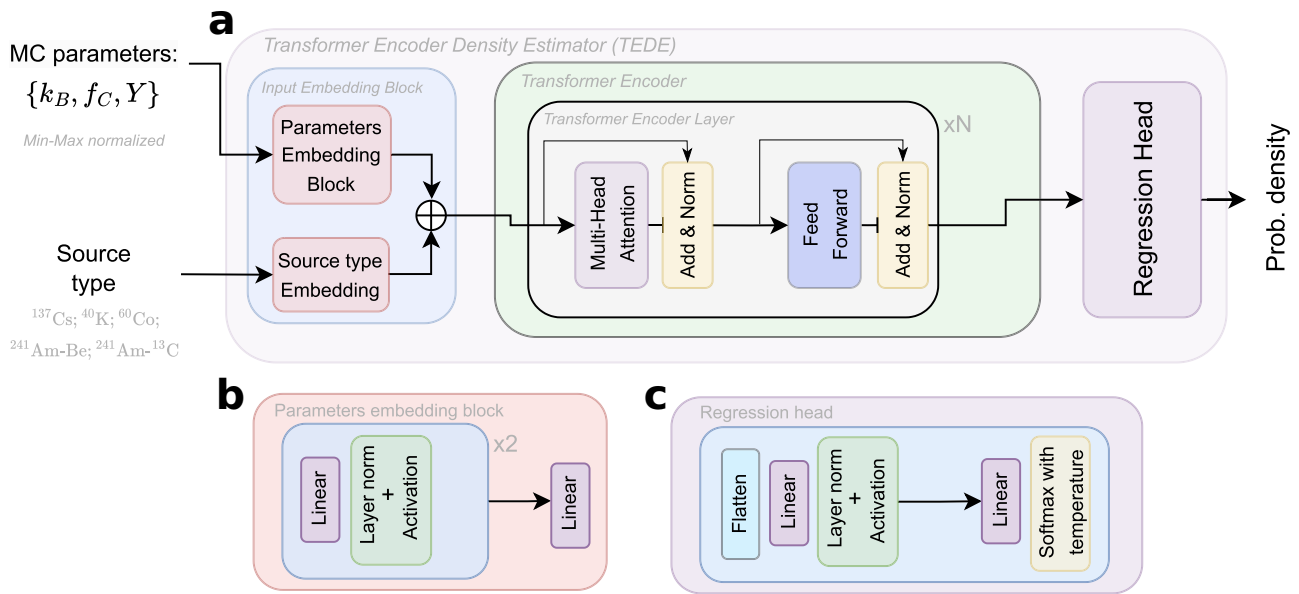


Fig. 8 | Schematic representation of the transformer model architecture. The Transformer Encoder Density Estimator (TEDE, **a**) takes min-max normalized energy response parameters (k_B, f_C, Y) and source type as inputs, processing them through dedicated embedding layers (**b**). The outputs of the embeddings are concatenated and fed into a transformer encoder comprising several encoder layers with multi-head attention. The transformer encoder’s output is then flattened and

processed by a regression head (**c**), with the softmax activation function as the final layer, that produces bin-by-bin probability density estimates for the corresponding energy spectrum. The transformer architecture is used as a powerful tool to capture complex relationships between the input energy response parameters and the calibration source probability densities.

transformer encoder architecture⁴⁶, which is widely used in the ML literature^{69,70} and becoming popular in particle physics applications in many different domains as well^{71–77}.

Architecture. The TEDE model processes inputs through three main components, as shown schematically in Fig. 8: (1) An input embedding block that transforms the input parameters ψ into high-dimensional vector representations. This block employs separate mechanisms for numerical and categorical inputs: the numerical MC parameters, after being min-max normalized to the range $[0, 1]$, are processed using a fully connected neural network with layer normalization⁷⁸, while the categorical source type is converted to a continuous vector representation through a trainable embedding layer. The resulting MC parameters and source type embeddings are then concatenated to form the input sequence for the next component, while each MC parameter is encoded with n_{tokens} —a tunable hyperparameter. Here, n_{tokens} controls the number of embedding vectors representing each continuous parameter in the transformer’s input sequence, affecting the model’s representational capacity. (2) A transformer encoder, the core of the model, consisting of several encoder layers (the exact number determined by hyperparameter optimization). Each layer utilizes multi-head self-attention mechanisms to capture complex dependencies among the embedded input features. Dropout regularization⁷⁹ can be applied within the encoder layers to mitigate overfitting; (3) A regression head, also constructed using a fully connected neural network with layer normalization, which processes the final representation from the transformer encoder. This head outputs a vector of $N_b = 800$ raw values (logits), $z = \{z_1, \dots, z_{N_b}\}$. These logits are then passed through the final temperature-scaled softmax layer and divided by the bin width Δx to create the output probability density estimates \hat{p}^{TEDE} :

$$\hat{p}_i^{\text{TEDE}} = \frac{1}{\Delta x} \frac{e^{z_i/T}}{\sum_{j=1}^{N_b} e^{z_j/T}}, \quad (8)$$

where z_i are the logits (outputs of the linear layer preceding the softmax) for bin i , and T is the softmax temperature hyperparameter. This construction inherently satisfies the PDF normalization requirement: $\sum_{i=1}^{N_b} \hat{p}_i^{\text{TEDE}} \Delta x = 1$.

Loss function. As a standard information-theoretic measure for comparing probability distributions, we use the Kullback-Leibler (KL) divergence^{80,81}, also known as relative entropy⁸², that quantifies the discrepancy between the target p^{JSW} and predicted \hat{p}^{TEDE} distributions. We define the KL divergence from \hat{p}^{TEDE} to p^{JSW} as follows:

$$\begin{aligned} D_{\text{KL}}(p^{\text{JSW}} \parallel \hat{p}^{\text{TEDE}}) &= \sum_{i=1}^{N_b} p_i^{\text{JSW}} \log \frac{p_i^{\text{JSW}}}{\hat{p}_i^{\text{TEDE}}} = \\ &= \sum_{i=1}^{N_b} p_i^{\text{JSW}} (\log p_i^{\text{JSW}} - \log \hat{p}_i^{\text{TEDE}}), \end{aligned} \quad (9)$$

always hold by construction of the softmax-based model’s output.

A fundamental property of the KL divergence is established by Gibbs’ inequality⁸³, which states that for any probability distributions P and Q :

$$D_{\text{KL}}(P \parallel Q) \geq 0 \quad (10)$$

This inequality guarantees that the KL divergence is always non-negative. Furthermore, the equality $D_{\text{KL}}(P \parallel Q) = 0$ holds if and only if the two distributions are identical. Gibbs’ inequality provides the theoretical justification for using KL divergence as a loss function: minimizing $D_{\text{KL}}(p^{\text{JSW}} \parallel \hat{p}^{\text{TEDE}})$ with respect to the TEDE parameters θ directly encourages the model’s output \hat{p}^{TEDE} to approximate the true distribution p^{JSW} .

This direct mapping approach, based on histogram approximation of the density function employed by the TEDE model, serves as a benchmark for the exact conditional PDF modeling approach utilized by a normalizing flows model.

Normalizing Flows Density Estimator (NFDE)

The method. This approach directly models the continuous conditional density $p(x|\psi)$ using the principle of normalizing flows⁴⁷. In our study, we consider conditional normalizing flows that learn an invertible and differentiable transformation f_θ , parameterized by learnable parameters θ and

conditioned on ψ . Since the main goal is conditional density estimation, we build inverted normalizing flows⁴⁹ that map the complex data distribution $p(x|\psi)$ to a simple base distribution $p_Z(z)$ (a standard normal $\mathcal{N}(0, 1)$) for which the probability density is known analytically. The density $p(x|\psi)$ estimates are then computed via the change of variables formula:

$$\widehat{p}(x|\psi) = p_Z(f_\theta(x|\psi)) \left| \frac{\partial f_\theta(x|\psi)}{\partial x} \right|, \tag{11}$$

where the Jacobian determinant term accounts for the change in volume induced by the transformation f_θ . This approach allows for exact likelihood evaluation and efficient sampling from the learned distribution.

Architecture. Our NFDE implementation constructs the transformation f_θ as a sequence of N_{flows} flows. We consider two simple types of transformations: planar flows and radial flows⁴⁶. The planar flow applies a simple, yet expressive, invertible transformation of the form $z' = z + u \tanh(wz + b)$, where z is the input vector from the previous flow (or the original data x for the first flow), and z' is the output vector. This transformation provides local, direction-specific deformations of the input space. The radial flow applies a transformation of the form $z' = z + \frac{\alpha\beta(z-\gamma)}{\alpha+|z-\gamma|}$ ⁴⁹, which creates global, symmetric transformations around a reference point γ . These parameters ($\{w, u, b\}$ for planar flows; $\{\alpha, \beta, \gamma\}$ for radial flows) for each flow are parametrized by a dedicated conditioning neural network, shown schematically in Supplementary Fig. 10. Similarly to the TEDE's input processing presented in Fig. 8, this network takes the parameters $\psi = \{k_B, f_C, Y, S\}$ as input and processes them through an input block. The numerical parameters (k_B, f_C, Y), after being min-max normalized to the range $[0, 1]$, are processed by a fully connected network, and the categorical source type (S) is mapped to an embedding vector. These representations are concatenated and fed through another fully connected network to produce the required flow parameters. This conditioning mechanism makes the overall transformation f_θ dependent on ψ , enabling the model to learn how the energy spectrum's shape changes across different parameter settings and source types. By composing a sequence of such conditional flows, the model can capture highly complex, multi-modal conditional densities.

Loss function. The NFDE model is trained by directly maximizing the likelihood of the training data, which is equivalent to minimizing the KL divergence between the model's likelihood $\widehat{p}(x|\psi)$ and the target distribution $p(x|\psi)$ ⁸⁴. Expanding the KL divergence gives:

$$\begin{aligned} D_{\text{KL}}[p(x|\psi) \parallel \widehat{p}(x|\psi)] &= \int p(x|\psi) \log \frac{p(x|\psi)}{\widehat{p}(x|\psi)} dx = \\ &= \int p(x|\psi) \log p(x|\psi) dx - \int p(x|\psi) \log \widehat{p}(x|\psi) dx = \\ &= \mathbb{E}_{p(x|\psi)}[\log p(x|\psi)] - \mathbb{E}_{p(x|\psi)}[\log \widehat{p}(x|\psi)]. \end{aligned} \tag{12}$$

The first term depends only on the true data distribution and is therefore independent of the model parameters θ . It can be treated as a constant with respect to optimization. The second term involves $\widehat{p}(x|\psi)$ and depends on θ through the model. As a result,

$$\begin{aligned} \arg \min_{\theta} D_{\text{KL}}[p(x|\psi) \parallel \widehat{p}(x|\psi)] &= \\ = \arg \max_{\theta} \mathbb{E}_{p(x|\psi)}[\log \widehat{p}(x|\psi)], \end{aligned} \tag{13}$$

which shows that minimizing the KL divergence is equivalent to maximizing the expected log-likelihood under the data distribution. In practice, since $p(x|\psi)$ is unknown but we have a dataset consisting of N samples $\{x_i, \psi_i\}_{i=1}^N$, the expectation can be estimated by the MC method using the empirical average:

$$\mathbb{E}_{p(x|\psi)}[\log \widehat{p}(x|\psi)] \approx \frac{1}{N} \sum_{i=1}^N \log \widehat{p}(x_i|\psi_i). \tag{14}$$

Thus, we define the loss function $L(\theta)$ as the negative log-likelihood:

$$\begin{aligned} L(\theta) &= -\frac{1}{N} \sum_{i=1}^N \log \widehat{p}(x_i|\psi_i) = \\ &= -\frac{1}{N} \sum_{i=1}^N \left(\log p_Z(f_\theta(x_i|\psi_i)) + \log \left| \det \frac{\partial f_\theta(x_i|\psi_i)}{\partial x_i} \right| \right). \end{aligned} \tag{15}$$

This formulation allows the model to learn the full conditional PDF without requiring discretization (binning) or imposing strong parametric assumptions on the distribution's shape, leveraging the unbinned nature of the input data.

Both models were implemented in PyTorch⁸⁵ and trained using PyTorch Lightning⁸⁶ framework. During training, we also employ different distributed strategies—GPU acceleration for TEDE and multi-CPU parallelization for NFDE—to optimize use of available computational resources.

Metrics to quantify model performance

To systematically evaluate model performance during training and hyperparameter optimization, we employ a family of statistical distances based on the L_p -norm framework to quantify differences between probability distributions. This approach straightforwardly unifies three classical statistical distances—Wasserstein, Cramér-von Mises, and Kolmogorov-Smirnov—within a single mathematical formalism. Specifically, for probability distributions U and V with respective cumulative distribution functions (CDFs) U and V , the L_p -norm distance is defined as:

$$d_p(u, v) = \left(\int_{-\infty}^{+\infty} |U(x) - V(x)|^p dx \right)^{1/p} \tag{16}$$

This formulation yields Wasserstein, Cramér-von Mises, and Kolmogorov-Smirnov distances depending on the choice of p :

- For $p = 1$, we obtain the Wasserstein distance, which measures the minimum “work” required to transform one distribution into another, making it particularly sensitive to shifts in distribution mean or location. Intuitively, it represents the minimum cost of transforming one distribution into another, where cost is proportional to the amount of probability mass moved multiplied by the distance it travels. It corresponds to the integrated absolute difference between the CDFs:

$$d_1(u, v) = \int_{-\infty}^{+\infty} |U(x) - V(x)| dx \tag{17}$$

- For $p = 2$, we obtain the Cramér-von Mises distance, which in the one-dimensional case is directly related to the energy distance (differing only by a factor of $\sqrt{2}$)⁸⁷. This metric balances sensitivity to both global distribution shifts and local deformations, effectively measuring the L_2 -norm of the difference between CDFs:

$$d_2(u, v) = \sqrt{\int_{-\infty}^{+\infty} |U(x) - V(x)|^2 dx} \tag{18}$$

- For $p \rightarrow \infty$, we obtain the Kolmogorov-Smirnov distance, which measures the maximum absolute difference between the two CDFs (L_∞ -norm) and is thus highly sensitive to localized discrepancies. This makes it particularly valuable for detecting localized anomalies that might be averaged out by other metrics:

$$d_\infty(u, v) = \sup_x |U(x) - V(x)| \tag{19}$$

This unified L_p -norm framework enables consistent

Table 3 | Comprehensive comparison of hyperparameter optimization for both models

Hyperparameters	TEDE	NFDE
Model architecture hyperparameters		
n_{tokens}	$\{k k \in \mathbb{N}, 1 \leq k \leq 5\}$: 1	---
d_{model}	$\{50k k \in \mathbb{N}, 1 \leq k \leq 10\}$: 100	---
n_{head}	$\{5, 10, 25\}$: 10	---
n_{layers}	$\{k k \in \mathbb{N}, 1 \leq k \leq 5\}$: 4	---
d_{ff}	$\{32k k \in \mathbb{N}, 1 \leq k \leq 15\}$: 384	---
n_{flows}	---	$\{100 + 5k k \in \mathbb{Z}, 0 \leq k \leq 20\}$: 120
n_{units}	---	$\{10 + 5k k \in \mathbb{Z}, 0 \leq k \leq 8\}$: 30
Flow type	---	Planar (fixed)
Dropout rate	$\{\rho_{\text{drop}} 0.0 \leq \rho_{\text{drop}} \leq 0.5\}$: 0.0	---
Activation function ^{91–94}	ReLU, GeLU	ReLU, GeLU, Tanh, SiLU
Softmax temperature	$\{T 0.25 \leq T \leq 3.0\}$: 2.03	---
General training hyperparameters		
Optimizer ^{95,96}	AdamW, RMSprop	AdamW (fixed)
Learning rate	$\{\eta 10^{-5} \leq \eta \leq 10^{-2}\}$: 1.9×10^{-3}	10^{-4} (fixed)
Scheduler type ^{97,98}	Exponential, ReduceOnPlateau, CosineAnnealing, None	ReduceOnPlateau, CosineAnnealing
Weight decay	$\{\lambda 10^{-6} \leq \lambda \leq 10^{-1}\}$: 8×10^{-5}	10^{-4} (fixed)
Batch size	$\{2^k k \in \mathbb{N}, 4 \leq k \leq 9\}$: 64	50 (fixed, 1 per CPU unit)
Optimizer-specific hyperparameters		
AdamW decay rate β_1	$\{\beta_1 0.5 \leq \beta_1 \leq 0.95\}$: 0.930	$\{\beta_1 0.5 \leq \beta_1 \leq 0.95\}$: 0.87
AdamW decay rate β_2	$\{\beta_2 0.9 \leq \beta_2 \leq 0.999\}$: 0.929	$\{\beta_2 0.9 \leq \beta_2 \leq 0.999\}$: 0.901
RMSprop parameter α	$\{\alpha 0.9 \leq \alpha \leq 0.999\}$: –	---
Scheduler-specific hyperparameters		
Cosine annealing period T_{max}	$\{T_{\text{max}} 5 \leq T_{\text{max}} \leq 75\}$: –	$\{T_{\text{max}} 5 \leq T_{\text{max}} \leq 75\}$: 50
Exponential rate τ	$\{\tau 0.8 \leq \tau \leq 0.99\}$: 0.92	---
Reduce on plateau factor γ	$\{\gamma 0.8 \leq \gamma \leq 0.99\}$: –	$\{\gamma 0.8 \leq \gamma \leq 0.99\}$: –
Optimization procedure setting		
Number of trials	250	25
Accelerator	GPU	50 parallel CPUs

The table presents search spaces and selected values (in **bold**) for both the Transformer Encoder Density Estimator (TEDE) and Normalizing Flows Density Estimator (NFDE). The table’s rows are organized into several subsections by hyperparameters application domain as follows: model architecture, general training, optimizer, learning rate scheduler, and the optimization procedure, respectively. For TEDE, the model architecture hyperparameters that define the model’s structure are: n_{tokens} (the number of tokens that encode each input parameter), d_{model} (the embedding dimension size of the transformer encoder), n_{head} (the number of attention heads), n_{layers} (the number of encoder layers), and d_{ff} (the encoder’s feed-forward dimension), dropout rate (controls the regularization strength), and the softmax temperature (the scaling factor for output probability distribution). For NFDE, the corresponding hyperparameters are: n_{flows} (the number of flows), n_{units} (the number of hidden units in each flow’s conditioning neural network), and flow type (the class of the transformations; kept fixed as Planar based on preliminary experiments). Common parameters include activation functions, optimizer type, and scheduler type (with their specific hyperparameters). For optimizer hyperparameters, AdamW uses β_1 and β_2 as exponential decay rates for moment estimates, while RMSprop uses α as the smoothing constant. For schedulers, τ represents the decay rate for Exponential and γ is the reduction factor for ReduceOnPlateau, while T_{max} is the half-cycle period for CosineAnnealing. Finally, the weight decay controls L_2 regularization strength, and batch size determines the number of different spectra processed simultaneously. Both models were optimized using Tree-structured Parzen Estimator, minimizing Cramér-von Mises distance (d_2) on the validation datasets.

evaluation across both TEDE and NFDE approaches. While these metrics naturally accommodate the continuous probability densities of NFDE, the binned case of TEDE can be accomplished with the integral being approximated by a sum. All three metrics were monitored throughout training and hyperparameter optimization, and we selected the Cramér-von Mises distance (d_2) as our primary optimization criterion due to its balanced sensitivity to both global distribution shifts and local deformations. This balance is particularly valuable when evaluating energy spectrum predictions across multiple calibration sources with diverse characteristics. We also found that d_2 distance provides the most robust translation to the ultimate goal of the models: energy response parameter estimation.

Hyperparameter optimization and model training

In the context of machine learning, the term “hyperparameters” refers to parameters that define a model’s architecture and its learning process. These parameters are set prior to the training process and must be optimized independently from the learnable parameters of the model (such as weights in a neural network).

Table 3 presents a comprehensive summary of the hyperparameter optimization spaces and results. For both models, we conduct systematic hyperparameter optimization using Optuna⁸⁸ with the Tree-structured Parzen Estimator (TPE)⁸⁹ algorithm. TPE is an efficient Bayesian optimization approach that sequentially builds probabilistic models of the hyperparameter-to-objective function relationship, focusing search efforts on promising regions of the hyperparameter space. As the objective function for the hyperparameter search, the average of the Cramér-von Mises

distances on the validation datasets is used:

$$d_2^{\text{Val}} = \frac{1}{2} \left[\frac{1}{N_{\text{Val}_1}} \sum_{i=1}^{N_{\text{Val}_1}} d_{2,i}^{\text{Val}_1} + \frac{1}{N_{\text{Val}_2}} \sum_{j=1}^{N_{\text{Val}_2}} d_{2,j}^{\text{Val}_2} \right], \quad (20)$$

where Val_1 is validation dataset 1 ($N_{\text{Val}_1} = 1000$ different parameter grid points) and Val_2 is validation dataset 2 ($N_{\text{Val}_2} = 3$ different parameter grid points). Thus, we select the set of hyperparameters that best minimizes d_2^{Val} .

In addition, we employ early stopping during each optimization trial, halting the training process when the model performance on validation data stops improving for a predefined number of epochs, the so-called ‘‘patience’’ (set to 200 for TEDE and 50 for NFDE).

For TEDE, we perform a comprehensive search across 250 trials, optimizing hyperparameters that can be split into two categories: (1) related directly to the model’s architecture and (2) related to the training procedure. The architecture hyperparameters are focused on values defining the transformer encoder structure⁴⁶: the number of tokens that encode each input MC parameter n_{tokens} , the embedding dimension d_{model} , the number of attention heads n_{heads} , the number of encoder layers n_{layers} , the feed-forward dimension d_{ff} , the dropout rate p_{drop} , the activation function type, and the softmax temperature. As for general training hyperparameters, we explore optimizer selection (AdamW vs. RMSprop), learning rate and its scheduling strategies, weight decay, and batch size. The batch size is defined as the number of different calibration source spectra processed simultaneously. Additionally, various optimizer- and scheduler-specific hyperparameters were tuned.

For NFDE, we perform 25 trials, focusing on flow-specific hyperparameters: the number of flows (n_{flows}) and the hidden unit dimension of the conditioning neural networks (n_{units}), which controls their capacity. To optimize computational resource costs for the hyperparameter search space, we first performed preliminary manual optimization using a trial-and-error approach to identify the most promising hyperparameters choices and fix them. Based on these preliminary experiments, planar flows consistently outperformed radial flows across different configurations, so the flow type was fixed as planar for the systematic optimization. Several other training hyperparameters (optimizer type, base learning rate, weight decay) were also fixed based on these preliminary experiments, while others, particularly scheduler-related parameters and the β_1, β_2 decay rates for the AdamW optimizer, underwent systematic optimization.

The best-found configurations (highlighted in bold in Table 3) show the following patterns: TEDE benefited from a moderate-sized transformer with no dropout ($p_{\text{drop}} = 0.0$) and with a softmax temperature more than twice the value of one ($T = 2.03$), that smooths the output probability distributions, and an Exponential learning rate scheduler. Meanwhile, NFDE performed best with 120 planar transformations, moderate conditioning network size ($n_{\text{units}} = 30$), and a CosineAnnealing learning rate scheduler. Both models favored the GeLU activation function, and both converged on similar optimization strategies with AdamW, although with different specific hyperparameter values ($\beta_1, \beta_2, T_{\text{max}}$) reflecting their approach-specific differences.

After the hyperparameter optimization procedure, both models were re-trained from scratch using the resulting optimized sets of the hyperparameters. The training dynamics, which illustrates the minimization of the average Cramér-von Mises distance for the validation data, is shown for both models in Supplementary Fig. 11. The early stopping patience was 200 and 100 epochs for TEDE and NFDE, respectively. Note that the final patience value for NFDE model was set to be twice higher than the cosine annealing period T_{max} to ensure proper convergence through multiple learning rate cycles.

Parameter estimation

We estimate the energy response parameters (k_B, f_C, Y) by comparing the ML model predictions to target data (from the testing datasets) using both

Bayesian and Frequentist inference. In the Frequentist framework, the `iminuit` package⁵⁹ is used for cost-function minimization, while in the Bayesian framework, we utilize a nested sampling method⁵¹ from `ultranest`⁵⁰ alongside a self-coded Metropolis-Hastings algorithm^{60,61} to sample the posterior distributions.

Likelihood functions for statistical inference. Let the parameter vector be denoted by $\phi = \{k_B, f_C, Y\}$, and the total dataset by \mathcal{D} . For a specific source, we define the corresponding subset of data as \mathcal{D}_s and the overall normalization constant N_s as the total number of expected events for that source. Here and throughout this subsection, the index s denotes a specific calibration source, taken from the following list: ¹³⁷Cs, ⁴⁰K, ⁶⁰Co, ²⁴¹Am-Be, ²⁴¹Am-¹³C.

For both models, the likelihood function used for statistical inference quantifies the probability of observing the data given the parameters, where the parameters to be optimized are ϕ and N_s .

The TEDE model (binned data). The TEDE model is designed for binned data, where the observation in each of the N_b bins is given by $m_{s,i}$ (for $i = 1, \dots, N_b$). The expected count of events for each bin can be found as $\mu_{s,i}(\phi) = N_s \cdot \hat{p}_i^{\text{TEDE}}(\phi, S)$. The likelihood based on the Poisson distribution for a source s is given by:

$$\mathcal{L}_s(\mathcal{D}_s | \phi) = \prod_{i=1}^{N_b} \frac{(\mu_{s,i}(\phi))^{m_{s,i}} e^{-\mu_{s,i}(\phi)}}{(m_{s,i})!}. \quad (21)$$

Taking the natural logarithm results in the log-likelihood:

$$\ln \mathcal{L}_s(\mathcal{D}_s | \phi) = \sum_{i=1}^{N_b} \left[m_{s,i} \cdot \ln \mu_{s,i}(\phi) - \mu_{s,i}(\phi) - \ln(m_{s,i}!) \right]. \quad (22)$$

Because the constant term $\ln(m_{s,i}!)$ does not depend on ϕ , it can be neglected during optimization. For the Bayesian method, this Poisson log-likelihood is used directly in constructing the posterior distribution. In the Frequentist approach, a related quantity known as the Poisson log-likelihood ratio is employed. Frequently, one minimizes the test statistic defined as:

$$\begin{aligned} -2\Delta \ln \mathcal{L}_s(\mathcal{D}_s | \phi) &= -2\ln \mathcal{L}_s(\mathcal{D}_s | \phi) + 2\ln \mathcal{L}_s(\mathcal{D}_s | \mathcal{D}_s) = \\ &= 2 \sum_{i=1}^{N_b} \left[\mu_{s,i}(\phi) - m_{s,i} + m_{s,i} \ln \frac{m_{s,i}}{\mu_{s,i}(\phi)} \right], \end{aligned} \quad (23)$$

where $\mathcal{L}_s(\mathcal{D}_s | \mathcal{D}_s)$ represents the likelihood of the saturated model in which the expected counts perfectly match the observed data ($\mu_{s,i} = m_{s,i}$), providing a reference for the maximum achievable likelihood. The convention $m_{s,i} \ln(m_{s,i}/\mu_{s,i}(\phi)) = 0$ when $m_{s,i} = 0$ is adopted. This test statistic asymptotically follows a χ^2 distribution and is particularly useful when comparing nested models or evaluating goodness-of-fit in a Frequentist framework. Compared to (22), it differs only by a data-dependent constant that does not affect minimization results.

The NFDE model (unbinned data). For NFDE, the analysis is performed on unbinned data. Let the individual data points for source S be denoted as $\{x_{s,i}\}_{i=1}^{M_s}$, where M_s is the total number of data points. The NFDE model outputs a PDF $\hat{p}(x_{s,i} | \phi, S)$.

With these ingredients, the extended unbinned likelihood^{57,58} can be constructed as a product of a Poisson term and terms involving the PDF:

$$\mathcal{L}_s(\mathcal{D}_s | \phi) = \frac{N_s^{M_s} e^{-N_s}}{M_s!} \cdot \prod_{i=1}^{M_s} \hat{p}(x_{s,i} | \phi, S). \quad (24)$$

Thus, the corresponding extended log-likelihood becomes:

$$\begin{aligned} \ln \mathcal{L}_s(\mathcal{D}_s | \phi) &= \\ &= M_s \ln N_s - N_s - \ln(M_s!) + \sum_{i=1}^{M_s} \ln \hat{p}(x_{s,i} | \phi, S), \end{aligned} \quad (25)$$

where $\ln(M_s!)$ term may be omitted during optimization since it does not depend on ϕ . This formulation naturally accounts for fluctuations in the total number of events due to the Poisson component, while the distribution shape is governed by the PDF term. The likelihood (25) is used directly in Bayesian analysis and scaled by a factor of (-2) for Frequentist analysis.

Extension to multiple spectra. Since full information originates from multiple independent sources, our framework is required to account for their combined contribution to parameter estimation. Using the five individual likelihood functions described earlier—one for each calibration source—the total likelihood function across all sources is defined as follows:

$$\mathcal{L}(\mathcal{D} | \phi) = \prod_{s=1}^5 \mathcal{L}_s(\mathcal{D}_s | \phi). \quad (26)$$

Similarly, since the logarithm transforms products into sums, the corresponding log-likelihood function can be expressed as:

$$\ln \mathcal{L}(\mathcal{D} | \phi) = \sum_{s=1}^5 \ln \mathcal{L}_s(\mathcal{D}_s | \phi). \quad (27)$$

We assume that data from different sources are statistically independent. Consequently, in this formulation each spectrum contributes independently to the total likelihood.

Inference approaches. After constructing the appropriate likelihood functions, the next step is to infer the parameter vector ϕ . In our study, two different approaches are employed:

Frequentist inference. For the Frequentist perspective, the parameter estimation is performed by maximizing the likelihood (or equivalently, minimizing the negative logarithm of the likelihood). The minimization is carried out with the `iminuit` package, which effectively finds the maximum likelihood estimates (MLEs) using the cost functions defined above.

Furthermore, the uncertainties in the parameter estimates are obtained using the MINOS algorithm implemented within `iminuit`. MINOS determines parameter uncertainties by scanning the profile of the likelihood function around the minimum, thereby accounting for potential asymmetries in the error estimates. Unlike simple symmetric error approximations (which are valid only when the likelihood surface is nearly quadratic), MINOS provides robust, asymmetric uncertainties, ensuring more reliable confidence intervals even when the likelihood is non-linear with respect to specific parameters⁹⁰.

Bayesian inference. In the Bayesian approach, the inferential goal is to obtain the posterior distribution of the parameters given the observed data, $q(\phi | \mathcal{D})$. Bayes' theorem states that:

$$q(\phi | \mathcal{D}) \propto \mathcal{L}(\mathcal{D} | \phi) \pi(\phi), \quad (28)$$

where $\pi(\phi)$ is the prior distribution encoding any previous knowledge about ϕ , which is considered flat in our study. The specific likelihood functions used are the ones outlined above.

To sample the posterior distribution, we utilize `UltraneSt`, a robust nested sampling algorithm, which effectively explores the parameter space even in complex high-dimensional settings. Nested sampling works by iteratively sampling from constrained prior volumes with progressively

higher likelihood thresholds, systematically shrinking the sampling region while maintaining uniform sampling within each likelihood contour. This approach enables efficient computation of both the posterior distribution and the Bayesian evidence (marginal likelihood) simultaneously. Moreover, a self-coded implementation of the Metropolis-Hastings algorithm is also applied to generate independent Markov chains, thereby providing complementary checks on the convergence and the robustness of the Bayesian inference.

The best-fit value of the parameters is determined as the median of the obtained posterior distribution, while the associated 1σ uncertainty is quantified using the 34.1% of the probability on either side of the median. This ensures a statistically robust characterization of the parameter estimates, avoiding biases that could arise from asymmetric posterior distributions.

Parameter estimation discussion. The dual approach, Frequentist optimization with `iminuit` (using the MINOS algorithm for robust uncertainty estimation) and Bayesian sampling with `ultranest` and Metropolis-Hastings, allows for cross-validation of parameter estimates while providing deep insights into the underlying uncertainties. The formulation of likelihood functions is inherently tied to the data structure: binned Poisson likelihoods for the TEDE model and extended unbinned likelihoods for the NFDE model ensure that the statistical features of the data are properly accounted for. Together, these methods yield not only point estimates from the MLE approach but also full posterior distributions via Bayesian inference, thus enhancing the overall robustness of our statistical conclusions.

Data availability

This study is based on simulation data produced within the JUNO collaboration. The data are not publicly available according to the collaboration data policy but can be shared upon request to the corresponding authors and with permission obtained from the collaboration.

Code availability

The code used for this study is available at <https://github.com/ArseniiGav/neuromct>.

Received: 4 August 2025; Accepted: 6 January 2026;

Published online: 23 January 2026

References

1. Capozzi, F. et al. Neutrino masses and mixing: entering the era of subpercent precision. *Phys. Rev. D.* **111**, 093006 (2025).
2. Agostini, M. et al. The Monte Carlo simulation of the Borexino detector. *Astropart. Phys.* **97**, 136–159 (2018).
3. An, F. P. et al. Measurement of the reactor antineutrino flux and spectrum at daya bay. *Phys. Rev. Lett.* **116**, 061801 (2016).
4. Wood, S. N. Statistical inference for noisy nonlinear ecological dynamic systems. *Nature* **466**, 1102–1104 (2010).
5. Papamakarios, G. & Murray, I. Fast ϵ -free inference of simulation models with Bayesian conditional density estimation. *NIPS* **29**, https://proceedings.neurips.cc/paper_files/paper/2016/hash/6aca97005c68f1206823815f66102863-Abstract.html (2016).
6. Lueckmann, J.-M. et al. Likelihood-free inference with emulator networks. *PMLR* **96**, 32–53 (2019).
7. Papamakarios, G., Sterratt, D. & Murray, I. Sequential neural likelihood: fast likelihood-free inference with autoregressive flows. *PMLR* **89**, 837–848 (2019).
8. Cranmer, K., Brehmer, J. & Louppe, G. The frontier of simulation-based inference. *Proc. Nat. Acad. Sci. USA* **117**, 30055–30062 (2020).
9. Lueckmann, J.-M. et al. Benchmarking simulation-based inference. *PMLR* **130**, 343–351 (2021).

10. Alsing, J., Wandelt, B. & Feeney, S. Massive optimal data compression and density estimation for scalable, likelihood-free inference in cosmology. *Mon. Not. Roy. Astron. Soc.* **477**, 2874–2885 (2018).
11. Alsing, J. et al. Fast likelihood-free cosmology with neural density estimators and active learning. *Mon. Not. Roy. Astron. Soc.* **488**, 4440–4458 (2019).
12. Kosiba, M. et al. The cosmological analysis of X-ray cluster surveys - VI. Inference based on analytically simulated observable diagrams. *Astron. Astrophys.* **693**, A46 (2025).
13. Brehmer, J. et al. Mining for Dark Matter Substructure: inferring subhalo population properties from strong lenses with machine learning. *Astrophys. J.* **886**, 49 (2019).
14. Dax, M. et al. Real-time gravitational wave science with neural posterior estimation. *Phys. Rev. Lett.* **127**, 241103 (2021).
15. Jeffrey, N. et al. Dark energy survey year 3 results: likelihood-free, simulation-based Λ CDM inference with neural compression of weak-lensing map statistics. *Mon. Not. Roy. Astron. Soc.* **536**, 1303–1322 (2024).
16. Brehmer, J. et al. MadMiner: machine learning-based inference for particle physics. *Comput. Softw. Big Sci.* **4**, 3 (2020).
17. Lazzarin, M., Alioli, S. & Carrazza, S. MCNNTUNES: tuning shower Monte Carlo generators with machine learning. *Comput. Phys. Commun.* **263**, 107908 (2021).
18. Heller, N. et al. Rejection sampling with autodifferentiation - case study: fitting a hadronization model 2411.02194 (2024).
19. Aad, G. et al. An implementation of neural simulation-based inference for parameter estimation in ATLAS. *Rept. Prog. Phys.* **88**, 067801 (2025).
20. Pina-Otey, S. et al. Likelihood-free inference of experimental neutrino oscillations using neural spline flows. *Phys. Rev. D.* **101**, 113001 (2020).
21. El Baz, M. & Sánchez, F. Fast posterior probability sampling with normalizing flows and its applicability in Bayesian analysis in particle physics. *Phys. Rev. D.* **109**, 032008 (2024).
22. Brehmer, J. et al. Mining gold from implicit models to improve likelihood-free inference. *Proc. Nat. Acad. Sci. USA* **117**, 5242–5249 (2020).
23. Brehmer, J. et al. Constraining effective field theories with machine learning. *Phys. Rev. Lett.* **121**, 111801 (2018).
24. An, F. et al. Neutrino physics with JUNO. *J. Phys. G* **43**, 030401 (2016).
25. Abusleme, A. et al. JUNO physics and detector. *Prog. Part. Nucl. Phys.* **123**, 103927 (2022).
26. Abusleme, A. et al. Potential to identify neutrino mass ordering with reactor antineutrinos at JUNO. *Chin. Phys. C.* **49**, 033104 (2025).
27. Abusleme, A. et al. Sub-percent precision measurement of neutrino oscillation parameters with JUNO. *Chin. Phys. C.* **46**, 123001 (2022).
28. Abusleme, A. et al. JUNO sensitivity to ${}^7\text{Be}$, pep, and CNO solar neutrinos. *JCAP* 022 (2023).
29. Zhao, J. et al. Model-independent approach of the JUNO ${}^8\text{B}$ solar neutrino program. *Astrophys. J.* **965**, 122 (2024).
30. Abusleme, A. et al. Real-time monitoring for the next core-collapse supernova in JUNO. *JCAP* 057 (2024).
31. Abusleme, A. et al. Prospects for detecting the diffuse supernova neutrino background with JUNO. *JCAP* 033 (2022).
32. Abusleme, A. et al. JUNO sensitivity to low energy atmospheric neutrino spectra. *Eur. Phys. J. C.* **81**, 10 (2021).
33. Abusleme, A. et al. JUNO sensitivity on proton decay $p \rightarrow \nu K^+$ searches. *Chin. Phys. C.* **47**, 113002 (2023).
34. Abusleme, A. et al. JUNO sensitivity to invisible decay modes of neutrons. *Eur. Phys. J. C.* **85**, 5 (2025).
35. Abusleme, A. et al. JUNO sensitivity to the annihilation of MeV dark matter in the galactic halo. *JCAP* 001 (2023).
36. Abusleme, A. et al. Prediction of energy resolution in the JUNO experiment. *Chin. Phys. C.* **49**, 013003 (2025).
37. Lin, T. et al. Simulation software of the JUNO experiment. *Eur. Phys. J. C.* **83**, 382 (2023).
38. Qian, Z. et al. Vertex and energy reconstruction in JUNO with machine learning methods. *Nucl. Instrum. Meth. A* **1010**, 165527 (2021).
39. Gavrikov, A., Malyshkin, Y. & Ratnikov, F. Energy reconstruction for large liquid scintillator detectors with machine learning techniques: aggregated features approach. *Eur. Phys. J. C.* **82**, 1021 (2022).
40. Yang, Z. et al. First attempt of directionality reconstruction for atmospheric neutrinos in a large homogeneous liquid scintillator detector. *Phys. Rev. D.* **109**, 052005 (2024).
41. Gavrikov, A. et al. Interpretable machine learning approach for electron antineutrino selection in a large liquid scintillator detector. *Phys. Lett. B* **860**, 139141 (2025).
42. Jiang, W. et al. Machine-learning based photon counting for PMT waveforms and its application to the improvement of the energy resolution in large liquid scintillator detectors. *Eur. Phys. J. C.* **85**, 69 (2025).
43. Malyshkin, Y. Novel fitting approach based on a neural network for JUNO. *EPJ Web Conf.* **337**, 01219 (2025).
44. Liu, J. et al. Neutrino type identification for atmospheric neutrinos in a large homogeneous liquid scintillation detector. *Phys. Rev. D.* **112**, 012018 (2025).
45. Abusleme, A. et al. Calibration strategy of the JUNO experiment. *JHEP* **03**, 004 (2021).
46. Vaswani, A. et al. Attention is all you need. *NIPS* **30**, https://papers.nips.cc/paper_files/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html (2017).
47. Tabak, E. G. & Turner, C. V. A family of nonparametric density estimation algorithms. *Commun. Pure Appl. Math.* **66**, 145–164 (2013).
48. Rezende, D. & Mohamed, S. Variational inference with normalizing flows. *PMLR* **37**, 1530–1538 (2015).
49. Trippe, B. L. & Turner, R. E. Conditional density estimation with Bayesian normalising flows 1802.04908 (2018).
50. Buchner, J. UltraNest - a robust, general purpose Bayesian inference engine. *J. Open Source Softw.* **6**, 3001 (2021).
51. Buchner, J. Nested sampling methods. *Stat. Surv.* **17**, 169–215 (2023).
52. Gavrikov, A. & Serafini, A. NeuroMCT: fast Monte Carlo tuning with generative machine learning in the JUNO experiment. *EPJ Web Conf.* **337**, 01119 (2025).
53. Goodfellow, I. J. et al. Generative adversarial networks. *NIPS* **27**, https://papers.nips.cc/paper_files/paper/2014/hash/f033ed80deb0234979a61f95710dbe25-Abstract.html (2014).
54. Gulrajani, I. et al. Improved training of Wasserstein GANs. *NIPS* **30**, https://papers.nips.cc/paper_files/paper/2017/hash/892c3b1c6dccc52936e27cbd0ff683d6-Abstract.html (2017).
55. Abi, B. et al. Deep Underground Neutrino Experiment (DUNE), far detector technical design report, volume I: Introduction to DUNE. *JINST* **15**, T08008 (2020).
56. Abe, K. et al. Hyper-Kamiokande design report 1805.04163 (2018).
57. Cowan, G. *Statistical Data Analysis* (Clarendon Press, 1998).
58. Barlow, R. Extended maximum likelihood. *Nucl. Instrum. Meth. A* **297**, 496–506 (1990).
59. Dembinski, H. et al. scikit-hep/iminuit. *Zenodo* <https://doi.org/10.5281/zenodo.3949207> (2020).
60. Hastings, W. K. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika* **57**, 97–109 (1970).
61. Metropolis, N. et al. Equation of state calculations by fast computing machines. *J. Chem. Phys.* **21**, 1087–1092 (1953).

62. Belle, V. & Papantonis, I. Principles and practice of explainable machine learning. *Front. Big Data* **4**, 688969 (2021).
63. HEP ML Community. A living review of machine learning for particle physics. <https://iml-wg.github.io/HEPML-LivingReview/> (2025).
64. Dirmeier, S., Albert, C. & Perez-Cruz, F. Simulation-based inference for high-dimensional data using surjective sequential neural likelihood estimation 2308.01054 (2023).
65. Gelman, A. et al. *Bayesian Data Analysis* 3rd edn (CRC Press, 2013).
66. Agostinelli, S. et al. Geant4 — a simulation toolkit. *Nucl. Instrum. Meth. A* **506**, 250–303 (2003).
67. Allison, J. et al. Recent developments in Geant4. *Nucl. Instrum. Meth. A* **835**, 186–225 (2016).
68. Huang, Y. et al. The flash ADC system and PMT waveform reconstruction for the daya bay experiment. *Nucl. Instrum. Meth. A* **895**, 48–55 (2018).
69. Lin, T. et al. A survey of transformers. *AI Open* **3**, 111–132 (2022).
70. Islam, S. et al. A comprehensive survey on applications of transformers for deep learning tasks. *Expert Syst. Appl.* **241**, 122666 (2024).
71. Butter, A. et al. Jet diffusion versus JetGPT — modern networks for the LHC. *SciPost Phys. Core* **8**, 026 (2025).
72. Gavrikov, A., García Pardiñas, J. & Garfagnini, A. DINAMO: dynamic and INterpretable anomaly MONitoring for large-scale particle physics experiments. *Mach. Learn. Sci. Tech.* **6**, 035050 (2025).
73. Spinner, J. et al. Lorentz-equivariant geometric algebra transformers for high-energy physics. *NeurIPS* **37**, https://proceedings.neurips.cc/paper_files/paper/2024/hash/277628cff838927d869cd1f671328ce0-Abstract-Conference.html (2024).
74. Wang, A. et al. Interpreting Transformers for Jet Tagging 2412.03673 (2024).
75. Pata, J. et al. Improved particle-flow event reconstruction with scalable neural networks for current and future particle detectors. *Commun. Phys.* **7**, 124 (2024).
76. Mokhtar, F. et al. Fine-tuning machine-learned particle-flow reconstruction for new detector geometries in future colliders. *Phys. Rev. D* **111**, 092015 (2025).
77. Dreyer, E. et al. Automated approach to accurate, precise, and fast detector simulation and reconstruction. *Phys. Rev. Lett.* **133**, 211902 (2024).
78. Ba, J. L., Kiros, J. R. & Hinton, G. E. Layer normalization 1607.06450 (2016).
79. Srivastava, N. et al. Dropout: a simple way to prevent neural networks from overfitting. *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
80. Kullback, S. & Leibler, R. A. On information and sufficiency. *Ann. Math. Stat.* **22**, 79–86 (1951).
81. Jordan, M. I. et al. An introduction to variational methods for graphical models. *Mach. Learn.* **37**, 183–233 (1999).
82. Dembo, A., Cover, T. M. & Thomas, J. A. Information theoretic inequalities. *IEEE Trans. Inf. Theory* **37**, 1501–1518 (1991).
83. Cover, T. M. & Thomas, J. A. *Elements of Information Theory* (Wiley-Interscience, 2006).
84. Papamakarios, G. et al. Normalizing flows for probabilistic modeling and inference. *J. Mach. Learn. Res.* **22**, 1–64 (2021).
85. Paszke, A. et al. Pytorch: an imperative style, high-performance deep learning library. *NeurIPS* **32**, https://papers.nips.cc/paper_files/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html (2019).
86. Falcon, W. et al. PyTorch lightning. <https://github.com/PyTorchLightning/pytorch-lightning> (2019).
87. Rizzo, M. L. & Székely, G. J. Energy distance. *Wiley Interdiscip. Rev. Comput. Stat.* **8**, 27–38 (2016).
88. Akiba, T. et al. Optuna: a next-generation hyperparameter optimization framework. In *Proc. ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (ACM)*, (2019).
89. Bergstra, J. et al. Algorithms for hyper-parameter optimization. *NIPS* **24**, https://papers.nips.cc/paper_files/paper/2011/hash/86e8f7ab32cfd12577bc2619bc635690-Abstract.html (2011).
90. Murtagh, B. A. & Saunders, M. A. Large-scale linearly constrained optimization. *Math. Program.* **14**, 41–72 (1978).
91. Nair, V. & Hinton, G. E. Rectified linear units improve restricted boltzmann machines. In *Proc. International Conference On Machine Learning (ICML)*, 807–814 <https://doi.org/10.5555/3104322.3104425> (2010).
92. Hendrycks, D. & Gimpel, K. Gaussian Error Linear Units (GELUs) 1606.08415 (2016).
93. Chen, F.-C. Back-propagation neural networks for nonlinear self-tuning adaptive control. *IEEE Control Syst. Mag.* **10**, 44–48 (1990).
94. Elfving, S., Uchibe, E. & Doya, K. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural Netw.* **107**, 3–11 (2018).
95. Loshchilov, I. & Hutter, F. Decoupled weight decay regularization. In *Proc. International Conference on Learning Representations* <https://openreview.net/forum?id=Bkg6RiCqY7> (2019).
96. Tieleman, T. & Hinton, G. Lecture 6.5-rmsprop: divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Netw. Mach. Learn.* **4**, 26–31 (2012).
97. Li, Z. & Sanjeev, A. An Exponential Learning Rate Schedule for Deep Learning. In *Proc. International Conference on Learning Representations* <https://openreview.net/forum?id=rJg8TeSFDH> (2020).
98. Loshchilov, I. & Hutter, F. SGDR: Stochastic gradient descent with warm restarts. In *Proc. International Conference on Learning Representations* <https://openreview.net/forum?id=Skq89Scxx> (2019).

Acknowledgements

We are thankful to the JUNO collaboration for the support and advice provided during the drafting of this manuscript. We are also very grateful to CNAF and JINR cloud services for providing the computing resources necessary for simulated data production and to CloudVeneto for offering IT support and infrastructure for training the machine learning models used in this study. Arsenii Gavrikov has received funding from the European Union's Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie Grant Agreement No. 101034319 and from the European Union — NextGenerationEU. Dmitry Dolzhikov and Maxim Gonchar are supported in the framework of the State project "Science" by the Ministry of Science and Higher Education of the Russian Federation under the contract 075-15-2024-541.

Author contributions

A.Gav., A.S., and M.Gr. conceived the study. A.Gav. generated simulation data, developed and trained the machine learning models. D.D., with support from A.S., developed the statistical inference framework and performed the parameter estimation analysis. A.Gav., A.S., and D.D. wrote the initial draft of the manuscript, with M.Go. and M.Gr. providing critical review of the draft. A.Gav. revised the manuscript. A.Gar., M.Go., and M.Gr. supervised the project. The remaining authors are members of the JUNO-Italia Consortium and contributed to the hardware and software development of the JUNO detector, making this study possible. All authors reviewed and approved the final manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-026-02499-6>.

Correspondence and requests for materials should be addressed to Arsenii Gavrikov, Andrea Serafini or Dmitry Dolzhikov.

Peer review information *Communications Physics* thanks Federico Sanchez Nieto and the other, anonymous, reviewer(s) for their contribution to the peer review of this work. [A peer review file is available].









































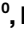


















Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2026

JUNO-Italia Consortium

Giuseppe Andronico ⁴, Vito Antonelli ⁵, Andrea Barresi ⁶, Davide Basilico ⁵, Marco Beretta ⁵, Antonio Bergnoli¹, Matteo Borghesi ⁶, Augusto Brigatti⁵, Riccardo Brugnera ¹, Riccardo Bruno⁴, Antonio Budano ⁷, Barbara Caccianiga ⁵, Antonio Cammi ⁸, Rossella Caruso ⁴, Vanessa Cerrone ¹, Davide Chiesa ⁶, Catia Clementi ⁹, Claudio Coletta ⁶, Lorenzo V. D'Auria ¹, Stefano Dusini ¹, Andrea Fabbri ⁷, Elia S. Farilla ⁷, Giulietto Felici ¹⁰, Giovanni Ferrante ⁶, Alberto Garfagnini ¹, Arsenii Gavrikov ¹ ✉, Marco G. Giammarchi ⁵, Nunzio Giudice⁴, Marco Grassi ¹, Nunzio Guardone⁴, Rosa Maria Guizzetti ¹, Fatima Houria ⁵, Cecilia Landini ⁵, Lorenzo Lastrucci ¹, Ivano Lippi ¹, Lorenzo Loi ⁸, Paolo Lombardi ⁵, Fabio Mantovani ^{11,12}, Stefano M. Mari ⁷, Agnese Martini ¹⁰, Lino Miramonti ⁵, Michele Montuschi ^{11,12}, Massimiliano Nastasi ⁶, Domizia Orestano ⁷, Fausto Ortica ⁹, Alessandro Paoloni ¹⁰, Luca Pelicci ⁵, Elisa Percalli ⁵, Fabrizio Petrucci ⁷, Ezio Previtali ⁶, Gioacchino Ranucci ⁵, Alessandra C. Re ⁵, Barbara Ricci ^{11,12}, Aldo Romani ⁹, Andrea Serafini ¹ ✉, Chiara Sirignano ¹, Monica Sisti ⁶, Luca Stanco ¹, Virginia Strati ^{11,12}, Marco D. C. Torri ⁵, Cristina Tuvè ⁴, Carlo Venettacci ⁷, Giuseppe Verde ⁴ & Lucia Votano ¹⁰

⁴INFN, Sezione di Catania e Università di Catania, Dipartimento di Fisica e Astronomia, Catania, Italy. ⁵INFN, Sezione di Milano e Università degli Studi di Milano, Dipartimento di Fisica, Milan, Italy. ⁶INFN, Sezione di Milano Bicocca e Dipartimento di Fisica Università di Milano Bicocca, Milan, Italy. ⁷INFN, Sezione di Roma Tre e Università degli Studi Roma Tre, Dipartimento di Matematica e Fisica, Rome, Italy. ⁸INFN, Sezione di Milano Bicocca e Dipartimento di Energetica, Politecnico di Milano, Milan, Italy. ⁹INFN, Sezione di Perugia e Università degli Studi di Perugia, Dipartimento di Chimica, Biologia e Biotecnologie, Perugia, Italy. ¹⁰Laboratori Nazionali dell'INFN di Frascati, Frascati, Italy. ¹¹INFN, Sezione di Ferrara, Ferrara, Italy. ¹²Università degli Studi di Ferrara, Dipartimento di Fisica e Scienze della Terra, Ferrara, Italy.