

Uncertainty-quantification-enabled inversion of nuclear responses

Krishnan Raghavan¹ and Alessandro Lovato^{2,3,4}

¹*Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, Illinois 60439, USA*

²*Physics Division, Argonne National Laboratory, Lemont, Illinois 60439, USA*

³*Computational Science Division, Argonne National Laboratory, Lemont, Illinois 60439, USA*

⁴*INFN-TIFPA Trento Institute of Fundamental Physics and Applications, Via Sommarive 14, 38123 Trento, Italy*



(Received 28 November 2023; revised 24 June 2024; accepted 2 August 2024; published 28 August 2024)

Nuclear quantum many-body methods rely on integral transform techniques to infer properties of electroweak response functions from ground-state expectation values. Retrieving the energy dependence of these responses is highly nontrivial, especially for quantum Monte Carlo methods, as it requires inverting the Laplace transform, a notoriously ill-posed problem. In this work, we propose an artificial neural network architecture suitable for accurate response function reconstruction with precise estimation of the uncertainty of the inversion. We demonstrate the capabilities of this new architecture benchmarking it against maximum entropy and previously developed neural network methods designed for a similar task, paying particular attention to its robustness against increasing noise in the input Euclidean responses.

DOI: [10.1103/PhysRevC.110.025504](https://doi.org/10.1103/PhysRevC.110.025504)

I. INTRODUCTION

The combination of sophisticated nuclear forces systematically derived within effective theories of QCD and numerical methods solving the quantum many-body problem with high accuracy [1–6] has enabled *abinitio* studies of the structure of several nuclides across the nuclear chart, including ²⁰⁸Pb [7]. Although existing many-body methods can describe nuclear ground-state properties and low-energy electroweak transitions with high accuracy [8], modeling real-time nuclear dynamics still poses a key challenge for current computational methods. Accessing it is crucial for achieving a fully microscopic understanding of processes such as fission, heavy-ion fusion, as well as lepton- and nucleus-nucleus scattering. Generally, computing dynamical properties of quantum many-body systems remains one of the paradigmatic open problems in quantum many-body theory, primarily due to quantum interference [9]. Computational limitations often strongly constrain the physical regimes in which quantum many-body dynamics can be solved on classical computers. Emerging technologies, such as neural-network quantum states [10,11] and quantum computing [12], hold great promise in this area, but their applications to nuclear physics are still in their infancy.

In this work, we will focus on the linear response regime, whose applications are ubiquitous in physics, including neutron scattering on materials and photon scattering in atomic systems and specifically on lepton-nucleus scattering. A quantitative description of the latter is critical for the interpretation of inclusive and semiexclusive electron-nucleus scattering experiments, shedding light on short-range correlations and the transition between hadronic and partonic degrees of freedom [13–15]. Additionally, the success of the accelerator neutrino program hinges on precise theoretical calculations of

neutrino-nucleus scattering, as they are essential for reconstructing the oscillated flux from measurements of particles produced in the aftermath of the scattering process [16–18].

State-of-the-art nuclear many-body methods, such as Green’s function Monte Carlo (GFMC) and coupled cluster, derive information about electroweak response functions from their integral transforms, which can be expressed as ground-state expectation values [19–21]; a notable exception in this regard consists in using an appropriate expansion in Chebyshev polynomials [22]. However, reconstructing the energy dependence of these response functions presents nontrivial challenges, especially when utilizing the Laplace kernel, as in the GFMC method [23].

The maximum entropy method (MaxEnt) [24,25], widely employed in condensed matter and lattice-QCD applications, has proved accurate in inverting the Laplace transform and reconstructing smooth response functions, characterized by a single broad quasielastic peak. On the other hand, MaxEnt struggles in the low-energy region, which is often characterized by several peaks, corresponding to low-energy nuclear transitions. For this reason, to retrieve the electromagnetic response of ¹²C these transitions had to be removed from the Euclidean response, using available experimental data [19]. Such shortcomings also yield certain discrepancies between GFMC and exact Faddeev results for the ³H muon capture rate near the nuclear breakup threshold, corresponding to energies of a few MeV [26].

Inspired by earlier machine-learning applications [27–31], in Ref. [32] a ‘physics informed’ artificial neural network (Phys-NN) was introduced for approximating the inverse of the Laplace transform. Phys-NN has been proved to outperform MaxEnt in both the low-energy transfer and the quasielastic regions, and to be more robust against noise in the input Euclidean responses. However, similarly to MaxEnt,

Phys-NN is not able to propagate the statistical uncertainties of the Euclidean response into the response function and to quantify the systematic errors due to the approximate inversion of the Laplace transform.

In this work, we overcome this limitation by developing an artificial neural network architecture that provides accurate response functions with quantified uncertainties, dubbed ‘‘UQ-NN.’’ To achieve this goal, we capitalize on a flexible parametrization of the response functions inspired by the one used in MaxEnt, which guarantees fast convergence of the training phase. As a result, UQ-NN exhibits an improved accuracy of the inversion and increased robustness to noise as compared to Phys-NN.

The present paper is structured as follows. In Sec. II we state the problem to be solved and discuss the relevant features of the nuclear electromagnetic responses. In Sec. III we describe our artificial-neural network architecture. In Sec. IV we present our results, and in Sec. V we discuss our conclusions.

II. ELECTROWEAK RESPONSES FROM THEIR LAPLACE TRANSFORMS

The nuclear response functions relevant to describing inclusive lepton-nucleus scattering cross sections in the one-boson exchange approximation can be generically written as

$$R(\mathbf{q}, \omega) = \int_{\mathcal{F}_f} \langle 0 | j^\dagger(\mathbf{q}, \omega) | f \rangle \langle f | j(\mathbf{q}, \omega) | 0 \rangle \times \delta(E_f - \omega - E_0). \quad (1)$$

In the above equation, $|0\rangle$ and $|f\rangle$ are the initial and final nuclear states, where the latter can be bound or unbound, with energies E_0 and E_f , respectively, and $j(\mathbf{q}, \omega)$ denotes the electroweak current operators.

In order to avoid computing all transitions induced by the current operator—which is impractical except for very light nuclear systems [33,34]—the GFMC method infers properties of the response functions from their Laplace transform [35], which is defined as

$$E_\alpha(\mathbf{q}, \tau) = \int_0^\infty d\omega e^{-\omega\tau} R_\alpha(\mathbf{q}, \omega). \quad (2)$$

Fixing the intrinsic energy dependence of the charge and current operators to the quasielastic (QE) peak, $\omega_{\text{QE}} = \sqrt{\mathbf{q}^2 + m^2} - m$, where m denotes the mass of the nucleon, one can express the Euclidean responses as ground-state expectation values:

$$E_\alpha(\mathbf{q}, \tau) = \langle 0 | j_\alpha^\dagger(\mathbf{q}, \omega_{\text{QE}}) e^{-(H-E_0)\tau} j(\mathbf{q}, \omega_{\text{QE}}) | 0 \rangle,$$

where H is the nuclear Hamiltonian. These expectation values can be evaluated by using the GFMC method on a uniform grid of n_τ imaginary-time points [23,35]. Standard GFMC calculations entail $n_\tau = 150$, with a maximum imaginary time of 0.07 MeV^{-1} . A set of noisy estimates for $E_\alpha(\mathbf{q}, \tau_i)$ can be obtained by performing independent imaginary-time propagations, from which the average Euclidean response $\bar{E}_\alpha(\mathbf{q}, \tau_i)$, and their associated statistical errors σ_i can be readily estimated.

In addition to the imaginary time, the continuous variables ω is also discretized on n_ω grid points, so that Eq. (2) becomes

$$E_i = \sum_{j=1}^{n_\omega} K_{ij} R_j \quad (3)$$

where $K_{ij} = e^{-\omega_j \tau_i} \Delta\omega_j$ and $R_j \equiv R(\mathbf{q}, \omega_j)$. For brevity, the dependence of the responses on the momentum transfer will be omitted from now on. The response functions reconstructed from GFMC calculations are typically tabulated on a grid of $n_\omega = 2000$ points with a maximum of $\omega_M = 2 \text{ GeV}$. The log-likelihood of the reconstructed responses is proportional to

$$\chi^2[\mathbf{R}, \bar{\mathbf{E}}] = \frac{1}{n_\tau} \sum_i \frac{(\sum_j K_{ij} R_j - \bar{E}_i)^2}{\sigma_i^2}. \quad (4)$$

where $\bar{\mathbf{E}} \in \mathbb{R}^{n_\tau}$, $\mathbf{R} \in \mathbb{R}^{n_\omega}$. Note that, GFMC calculations provide the sample covariance matrix between the data at $\tau = \tau_i$ and $\tau = \tau_j$, which is typically nondiagonal because of correlations among the imaginary-time points [19,23,35]. However, the likelihood reduces to the one of Eq. (4) once the data and the Laplace kernel are rotated in the basis where the covariance is a diagonal [25].

Due to the smoothing effect of the Laplace kernel, a simple minimization of χ^2 results in multiple response functions that are consistent, within errors, with the GFMC Euclidean response. Maximum entropy methods [24,25] aim to address this ambiguity by treating the response functions, both positive definite and normalizable, as probability distributions. According to the principle of maximum entropy, their values are determined by maximizing the entropy, defined as

$$S[\mathbf{R}, \mathbf{M}] = \sum_{i=1}^{n_\omega} \left[R_i - M_i - R_i \ln \left(\frac{R_i}{M_i} \right) \right] \Delta\omega_i, \quad (5)$$

The positive valued function $M(\omega)$ serves as the default model and incorporates prior knowledge about $R(\omega)$ when no data is available and $\mathbf{M} \in \mathbb{R}^{n_\omega}$. The entropy quantifies the deviation between the response function and the default model and it is closely related with the Kullback-Leibler divergence [36]. It has a maximum value of zero when $\mathbf{R} = \mathbf{M}$ and it is negative otherwise. When reconstructing the response functions from GFMC calculations, a flat noninformative default model is typically used: $M(\omega) = \theta(\omega_M - \omega)$, where $\theta(\omega)$ is the step function.

By applying Bayes’ theorem, MaxEnt identifies the most probable response function, minimizing the quantity

$$Q[\mathbf{R}, \bar{\mathbf{E}}, \mathbf{M}] = \frac{1}{2} \chi^2[\mathbf{R}, \bar{\mathbf{E}}] - \alpha S[\mathbf{R}, \mathbf{M}], \quad (6)$$

with respect to \mathbf{R} . Here α is a parameter that determines the balance between the entropy and the log-likelihood. When $\alpha = 0$, the standard ill-posed minimization is recovered, and for $\alpha \gg 1$, \mathbf{R} converges to the default \mathbf{M} . In our study, we carry out all benchmarks against the historic MaxEnt approach [37], which selects α such that $\chi^2 = 1$. More sophisticated methods like the *classic* MaxEnt [38] and *Bryan* MaxEnt [24] tend to overfit the data [39,40]. The arbitrariness in choosing α hampers a robust reconstruction of $R(\omega)$ in the small- ω region. Specifically, too small α results in

overfitting $E(\tau)$ and uncontrolled oscillations in the reconstructed responses.

The key point in the inversion of the Laplace transform resides in the minimization of $Q[\mathbf{R}, \bar{\mathbf{E}}, \mathbf{M}]$, defined in Eq. (6) for given $\bar{\mathbf{E}}$ and \mathbf{M} . An efficient way to accomplish this task was first discussed in Ref. [24], and it entails performing a singular value decomposition (SVD) of the kernel

$$K = V \Sigma U^T. \quad (7)$$

In the above equation, U and V are $n_\omega \times n_\omega$ and $n_\tau \times n_\tau$ orthogonal matrices, while Σ is a $n_\tau \times n_\omega$ rectangular diagonal matrix. Since the kernel is effectively singular, the smallest elements on the diagonal are essentially zero for the numerical precision. Hence, without loss of accuracy, we keep only the n_s largest eigenvalues and disregard the others so that only the first n_s columns of U are relevant for representing the kernel.

The gradient of the log-likelihood is given by

$$\frac{\partial \chi^2}{\partial R_i} = \sum_j \frac{\partial \chi^2}{\partial E_j} \frac{\partial E_j}{\partial R_i} = K_{ij}^T \frac{\partial \chi^2}{\partial E_j}. \quad (8)$$

Since the columns of K^T are linear combinations of the ones of U , all the search directions for the minimum are spanned, within machine precision, by the first n_s columns of U . In this *singular space*, the stationary condition of $Q[R]$ reads

$$0 = \frac{\partial Q}{\partial R_i} = \alpha \frac{\partial S}{\partial R_i} - \frac{1}{2} \frac{\partial \chi^2}{\partial R_i} = 0, \quad (9)$$

which implies

$$-\alpha \ln(R_i/M_i) = \frac{1}{2} \sum_j K_{ij}^T \frac{\partial \chi^2}{\partial E_j}. \quad (10)$$

Thus, the solution can be represented in terms of a vector \mathbf{u}

$$\ln \left(\frac{R_i}{M_i} \right) = K_{ij}^T u_j. \quad (11)$$

Since only the first n_s elements of Σ are different from zero, not all the n_τ components of \mathbf{u} are independent. Since K^T and U share the same vector space and since most of the relevant search directions lie in the singular space, the solution can be written in the form

$$R_i = M_i \exp \left(\sum_{j=1}^{n_s} U_{ij} u_j \right). \quad (12)$$

Therefore, to the machine-precision level, the most general solution of Eq. (10) only depends on the n_s coordinates u_j . In MaxEnt applications, owing to the ranges of ω and τ in GFMC calculations, $n_s \simeq 30 \ll n_\omega = 2000$. Hence, a standard Newton procedure to minimize $Q[\mathbf{R}, \bar{\mathbf{E}}, \mathbf{M}]$ converges much faster for finding u_j than for the original R_i .

III. ARTIFICIAL NEURAL NETWORK INVERSION ALGORITHM

In a recent work [32], a physics-informed artificial neural network (Phys-NN) was introduced to approximate the inverse of the Laplace transform. Phys-NN employs a Gaussian kernel basis to capture the structure of the Laplace

kernel. In this work we utilize instead the more advantageous parametrization of Eq. (12), using artificial neural networks to determine the coefficients u_j . Formally, the reconstructed response is given by

$$\hat{R}_i(\boldsymbol{\theta}) = m_i \exp \left(\sum_{j=1}^{n_s} U_{ij} u_j(\boldsymbol{\theta}) \right), \quad (13)$$

where $\boldsymbol{\theta}$ denotes the collection of training parameters. This entails a critical reduction of the artificial neural network output dimension compared to Phys-NN, whose outputs were directly the n_ω values \hat{R}_i .

A. Entropy neural network (Ent-NN)

The first architecture discussed in this work, dubbed ‘‘Ent-NN,’’ takes as input the n_τ discrete Euclidean response values and provides the corresponding functions. The architecture of Ent-NN, displayed in Fig. 1, is comprised of three central elements: (i) the attention mechanism comprised of two feed-forward layers with one skip connection that takes as input E_i and generates the coefficients $u_i(\boldsymbol{\theta})$, (ii) the fixed basis function U_{ij} , used to estimate $\hat{\mathbf{R}}(\boldsymbol{\theta})$ from $u_i(\boldsymbol{\theta})$ as in Eq. (13), and (iii) the discrete Laplace transform of Eq. (3) for computing the Euclidean $\hat{E}(\boldsymbol{\theta})$ associated with the reconstructed response function as $\hat{E}_i(\boldsymbol{\theta}) = \sum_{j=1}^{n_\omega} K_{ij} \hat{R}_j(\boldsymbol{\theta})$.

1. Training

As in Ref. [32], Ent-NN is trained on two distinct datasets comprising pairs of physically relevant $R(\omega)$, $E(\tau)$. The responses belonging to the first dataset are characterized by a single broad asymmetric peak, corresponding to the QE reaction mechanism, modeled by a skew-normal distribution. The responses belonging to the second dataset exhibit a sharper elastic (EL) peak at low energy, in addition to the QE one. The corresponding Euclidean responses are obtained by applying the discrete Laplace transform of Eq. (3). Since the simulated responses are smooth functions of ω , the numerical integration error on the Euclidean responses is about 10^{-5} . For each of the one-peak and two-peaks cases, we generate a total of 500 000 pairs $(\mathbf{R}_k, \mathbf{E}_k) \in \mathbb{R}^{n_\omega+n_\tau}$ of responses and corresponding Euclidean responses, which we then partition into training (\mathbb{T}), validation (\mathbb{V}), and test/out-of-sample (\mathbb{O}) datasets. The one-peak and two-peak test datasets comprise 1000 pairs each; the combined test dataset is just the union of these two sets. We use 80% and 20% of the remaining data for training the network and validation, respectively.

The optimal values for the parameters $\boldsymbol{\theta}$ are found by the standard supervised learning approach of approximately solving

$$\min_{\boldsymbol{\theta}} \frac{1}{|\mathbb{T}|} \sum_{k \in \mathbb{T}} \ell[\mathbf{E}_k, \mathbf{R}_k, \hat{\mathbf{R}}_k(\boldsymbol{\theta})] \quad (14)$$

by using a minibatch-based stochastic gradient descent procedure to minimize an empirical loss function. Our overall objective in the above equation is the average loss over the $|\mathbb{T}|$ points in the training set. Taking inspiration from MaxEnt, for each data and model output, we employ a loss function that is

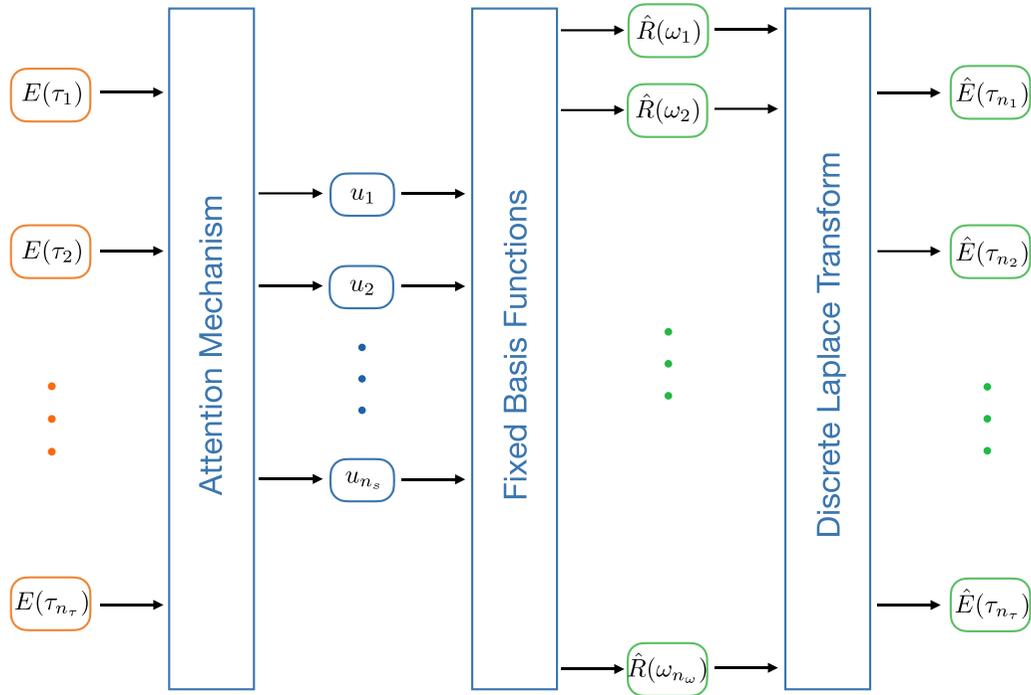


FIG. 1. Schematic overview of the Ent-NN architecture. The attention mechanism takes as input $E(\tau_i)$ and generates the coefficients $u_i(\theta)$. The fixed basis function uses the latter to output $\hat{R}_i = M_i \exp[\sum_{j=1}^{n_s} U_{ij} u_j(\theta)]$. The discrete Laplace transform numerically integrates the reconstructed response to obtain the corresponding Euclidean response as $\hat{E}_i(\theta) = \sum_{j=1}^{n_\omega} K_{ij} \hat{R}_j(\theta)$.

the sum of a *response* and a *Euclidean cost*

$$\ell[\mathbf{E}_k, \mathbf{R}_k, \hat{\mathbf{R}}_k(\theta)] = \gamma_R |S[\hat{\mathbf{R}}_k(\theta), \mathbf{R}_k]| + \gamma_E \chi^2[\hat{\mathbf{R}}_k(\theta), \mathbf{E}_k].$$

As discussed below, the positive-definite constants γ_R and γ_E are chosen to compensate for the fact that $\chi^2[\hat{\mathbf{R}}_k(\theta), \mathbf{E}_k]$ is typically much larger than the entropy $S[\hat{\mathbf{R}}_k(\theta), \mathbf{R}_k]$. The response cost—closely related to the Kullback-Leibler divergence [36]—ensures that the reconstructed response functions are close to the original ones. The absolute value ensures that the response cost has a minimum value of zero when $\mathbf{R}_k = \hat{\mathbf{R}}_k(\theta)$ and is positive otherwise. The Euclidean cost is aimed at aligning the Laplace transform $\hat{\mathbf{E}}(\theta)$ of the reconstructed response functions with the original Euclidean responses.

Since the inversion of the Laplace transform is an ill-posed problem, there are many response functions, possibly wildly different among each other, whose Laplace transform are compatible with the original Euclidean responses within statistical uncertainties. Consequently, there are instances in which χ^2 is small even when the reconstructed response is not similar to the original one, leading to potential instabilities in the minimization procedure. To tame this behavior, we split the training into two phases.

In the first phase, we take $\gamma_R = 10^7$ and $\gamma_E = 10^{-7}$ and optimize the network using the Adam [41] optimizer with a learning rate of 10^{-4} . Since $\gamma_R \gg \gamma_E$, the entropy response cost dominates the loss function and drives the reconstructed response functions close to the original ones. Once the entropy cost has reduced significantly, we enter the second phase of the optimization, where we keep $\gamma_R = 10^7$ while increasing the relative importance of the Euclidean cost by taking $\gamma_E = 1$

so that Ent-NN also learns to keep the Laplace transform of the response function close to the original Euclidean response. Reducing the learning rate in the second phase to 10^{-5} is necessary in order to keep the reconstructed response functions close to the optimal ones found in the previous phase.

To estimate the epistemic uncertainty of Ent-NN, we train ten neural networks using identical architectures but with independently initialized parameters. This approach, also employed in Ref. [32], enables us to assess how variations in the initial conditions affect the predictions of Ent-NN, providing insights into the accuracy and reliability of the model when applied to unseen data.

B. Uncertainty quantification neural network (UQ-NN)

Meaningful comparisons between GFMC calculations of the response function with experimental data require carrying out rigorous uncertainty quantification. The latter is particularly relevant when making predictions for neutrino-nucleus scattering, as cross-section uncertainties should be carefully propagated in the error budget of neutrino-oscillation parameters [42].

Approximately inverting the Laplace transform using artificial neural-network entails two distinct sources of uncertainty. The first one, discussed in the previous section, is due to the choice of the neural network model, which includes the set of optimal parameters found in the training procedure. In Ref. [32], this effect was found to be small. The second one concerns propagating the statistical errors associated with GFMC estimates of \mathbf{E} into the reconstructed responses.

The distribution of the computed Euclidean response computed within the GFMC is Gaussian, hence

$$P(\mathbf{E}) \propto \exp\left(-\frac{(\mathbf{E} - \bar{\mathbf{E}})^2}{2\sigma^2}\right). \quad (15)$$

Consistent with the notation of Sec. II, we assume a diagonal covariance matrix and, to simplify the discussion, we also consider the standard deviation σ to be independent of τ ; both assumptions can be easily relaxed. The corresponding probability distribution of the response functions is then given by

$$P(\mathbf{R}) = \int d\mathbf{E} P(\mathbf{R}|\mathbf{E})P(\mathbf{E}) \quad (16)$$

Assuming that the response functions can be accurately reconstructed using artificial neural networks and that the training parameters are narrowly distributed around the optimal ones implies that

$$P(\mathbf{R}|\mathbf{E}) = \delta(\mathbf{R} - \hat{\mathbf{R}}(\boldsymbol{\theta})). \quad (17)$$

Hence, following standard Monte Carlo error propagation procedures [43], samples of response functions distributed according to $P(\mathbf{R})$ of Eq. (16) can, in principle, be obtained by drawing Euclidean responses from the Gaussian distribution of Eq. (15) and applying Ent-NN to each of them.

The first task is accomplished by adding stochastic noise [31] to the simulated Euclidean responses as

$$E_i^\sigma = \bar{E}_i + \epsilon_i, \quad (18)$$

where ϵ_i are independent samples from a Gaussian distribution with mean zero and standard deviation σ [32]. This particular choice is consistent with GFMC calculations of Euclidean response functions, which follow a Gaussian distribution. We note that including an energy-dependent error σ_i , which would more closely resemble actual GFMC estimates, would involve drawing ϵ_i from Gaussian distributions characterized by energy-dependent widths.

As for the second step above, applying Ent-NN to the Euclidean obtained as in Eq. (18) results in exceedingly large variations in the reconstructed responses. The reason for this behavior has to be ascribed to the fact that adding Gaussian noise to the Euclidean responses makes them significantly different from the ones found in the original training dataset. Hence, Ent-NN is forced to extrapolate, leading to inaccurate reconstructions, as apparent by the significant discrepancies between the reconstructed Euclidean responses $\hat{\mathbf{E}}(\boldsymbol{\theta})$ and the noisy ones \mathbf{E}^σ that are taken as input.

To remedy this shortcoming, we have developed UQ-NN. Its architecture, schematically illustrated in Fig. 2, comprises two Ent-NN neural networks operating in parallel. The first takes as input the original \mathbf{E} and returns the corresponding $\hat{\mathbf{R}}(\boldsymbol{\theta})$ as well as $\hat{E}_i(\boldsymbol{\theta}) = \sum_{j=1}^{n_\omega} K_{ij} \hat{R}_j(\boldsymbol{\theta})$. The second takes as input the noisy \mathbf{E}^σ and returns $\hat{\mathbf{R}}^\sigma(\boldsymbol{\theta})$ and $\hat{E}_i^\sigma(\boldsymbol{\theta}) = \sum_{j=1}^{n_\omega} K_{ij} \hat{R}_j^\sigma(\boldsymbol{\theta})$. Crucially, the training parameters are common to the two Ent-NN blocks.

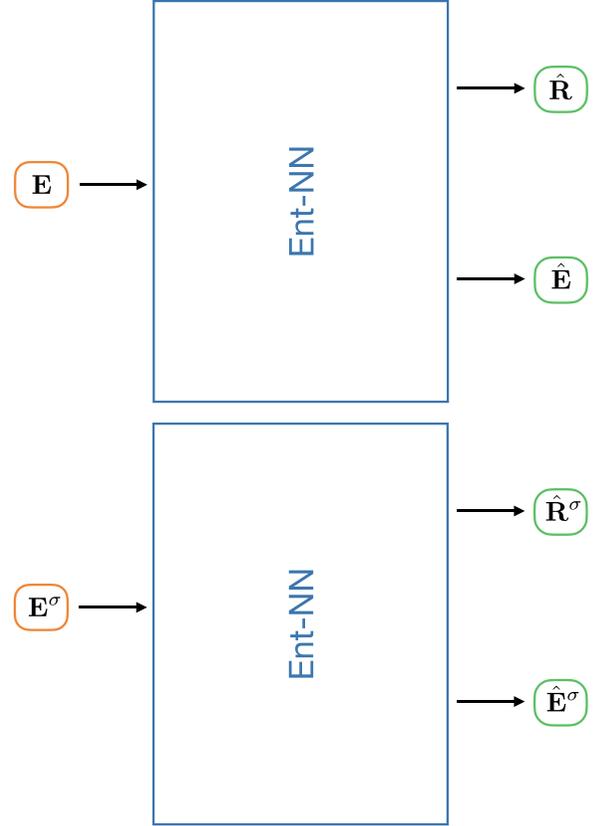


FIG. 2. Schematic overview of the UQ-NN architecture. Each of the two blocks denotes a single Ent-NN, whose outputs are $\hat{\mathbf{R}}$ and \hat{E} , with the latter obtained by computing the Laplace transform of the former numerically. Note that the two Ent-NN blocks are identical.

1. Training

The training process for UQ-NN follows the steps as in Ent-NN, with the distinction that each batch of data also includes the noisy Euclidean responses of Eq. (18). Hence, the optimal values of $\boldsymbol{\theta}$ are found by

$$\min_{\boldsymbol{\theta}} \frac{1}{|\mathbb{T}|} \sum_{k \in \mathbb{T}} \ell[\mathbf{E}_k, \mathbf{R}_k, \mathbf{E}_k^\sigma, \hat{\mathbf{R}}_k(\boldsymbol{\theta}), \hat{\mathbf{R}}_k^\sigma(\boldsymbol{\theta})] \quad (19)$$

Crucially, no noisy response appears among the arguments of the loss function, as there is no direct way to generate it from \mathbf{E}_k^σ . The loss function is a generalization of the one of Ent-NN:

$$\begin{aligned} \ell[\mathbf{E}_k, \mathbf{R}_k, \mathbf{E}_k^\sigma, \hat{\mathbf{R}}_k(\boldsymbol{\theta}), \hat{\mathbf{R}}_k^\sigma(\boldsymbol{\theta})] \\ = \gamma_R |S[\hat{\mathbf{R}}_k(\boldsymbol{\theta}), \mathbf{R}_k]| + \gamma_E \{\chi^2[\hat{\mathbf{R}}_k(\boldsymbol{\theta}), \mathbf{E}_k] + \chi^2[\hat{\mathbf{R}}_k^\sigma(\boldsymbol{\theta}), \mathbf{E}_k^\sigma]\}. \end{aligned} \quad (20)$$

The additional term in the Euclidean cost drives the reconstructed $\mathbf{E}^\sigma(\boldsymbol{\theta})$ to be close to the input noisy Euclidean, thereby providing a reliable reconstruction of the corresponding noisy response function. Specifically, for each Euclidean response \mathbf{E}_k in a batch of data, we also sample \mathbf{E}_k^σ as in Eq. (18). We then employ UQ-NN to evaluate the corresponding $\hat{\mathbf{R}}_k(\boldsymbol{\theta})$ and $\hat{\mathbf{R}}_k^\sigma(\boldsymbol{\theta})$. Finally, we numerically integrate them to obtain the corresponding Euclidean, which are used

in the cost function defined in Eq. (20). The rest of the hyperparameters and the choice of γ_R and γ_E are identical to the training of Ent-NN.

The training is performed for different noise levels. The smallest is $\sigma = 10^{-5}$, which is similar to the statistical noise of actual GFMC calculations of ${}^4\text{He}$ Euclidean electromagnetic responses. We also consider $\sigma = 10^{-4}$, $\sigma = 10^{-3}$ (a value compatible with typical GFMC calculations of ${}^{12}\text{C}$), and $\sigma = 10^{-2}$. This latter noise level corresponds to typical auxiliary-field diffusion Monte Carlo [44] (AFDMC) calculations of ${}^{16}\text{O}$ that are currently being performed.

IV. RESULTS

A. Model performance

To quantify the accuracy of both Ent-NN and UQ-NN, we adopt two metrics averaged over the test/out-of-sample dataset \mathbb{O} . The first one is the average absolute value of the entropy

$$\bar{S} = \frac{1}{|\mathbb{O}|} \sum_{k \in \mathbb{O}} |S[\hat{\mathbf{R}}_k(\boldsymbol{\theta}), \mathbf{R}_k]|,$$

where the entropy functional is defined in Eq. (5). It is important to note that a smaller \bar{S} corresponds to more accurate reconstructed responses. The second one is the average reduced χ^2_E ,

$$\overline{\chi^2_E} = \frac{1}{|\mathbb{O}|} \sum_{k \in \mathbb{O}} \chi^2[\hat{\mathbf{R}}_k(\boldsymbol{\theta}), \mathbf{E}_k],$$

where the log-likelihood is the one of Eq. (4). We note that, due to the smoothing nature of the Gaussian kernel, the entropy is the best-suited metric to compare the performances of different methods employed to invert the Laplace transform. Indeed, even responses that are wildly different from the original one might yield $\chi^2 = 1$. These responses would, in turn, produce very large S .

Table I summarizes the testing metrics for the single-peak, two-peak, and combined datasets, comparing Ent-NN and UQ-NN against Phys-NN and MaxEnt. Note that, while the MaxEnt metrics are identical to those found in Ref. [32], Phys-NN exhibits lower entropy and a larger χ^2 . This change in performance is a deliberate choice in the training process. To better reconstruct the responses, we found it beneficial to use a larger γ_E and a smaller γ_R in the loss function. In this initial comparison, no noise has been added to the input Euclidean responses, which only suffer from the numerical integration error discussed in Sec. III. All approaches perform best in reconstructing one-peak responses, while the accuracy of two-peak reconstructions appears to be lower. The reconstructions for the combined dataset fall between those of the other two datasets. This behavior is expected, considering that response functions characterized by two peaks, especially with the EL peak located in the low- ω region, are notoriously more complicated to reconstruct than those with a single broad QE peak.

Regarding the entropy metric, Ent-NN, UQ-NN, and Phys-NN significantly outperform MaxEnt across the one-peak, two-peak, and combined datasets. The improved performance

TABLE I. Ent-NN, UQ-NN, Phys-NN, and MaxEnt testing metrics \bar{S} , and $\overline{\chi^2}$ for the one-peak, two-peak, and combined datasets. The Ent-NN results are obtained by averaging the ten randomly initialized networks.

	$\overline{\chi^2}$	$\bar{S} \times 10^{-4}$
	Ent-NN	
One-peak	2.886	0.474
Two-peak	8.440	1.361
Combined	3.445	0.832
	UQ-NN	
One-peak	2.595	0.439
Two-peak	9.251	1.375
Combined	5.923	0.907
	Phys-NN	
One-peak	2.002	6.222
Two-peak	7.766	14.42
Combined	5.153	1.031
	MaxEnt	
One-peak	1.015	60.4
Two-peak	1.016	107
Combined	1.015	83.7

of Ent-NN and UQ-NN in capturing the energy dependence of the response functions compared to Phys-NN is a direct result of utilizing the basis functions outlined in Eq. (13), which are tailored for inverting the Laplace transform. Conversely, historic MaxEnt yields the least accurate reconstructions. Ent-NN and Phys-NN generate similar reconstruction values, but Ent-NN has a slight edge in terms of entropy. Although UQ-NN's reconstructions are less precise than Ent-NN's, they outperform both Phys-NN and MaxEnt.

In comparing the χ^2 values, MaxEnt appears to perform the best. However, this is due to the historical MaxEnt method, where the optimal response function is determined by setting α in Eq. (6) to achieve $\chi^2 = 1$. As indicated by the entropy metric, the inherently ill-posed nature of the problem implies that a $\chi^2 \approx 1$ does not guarantee an accurate reconstruction of the original response functions. In fact, despite Ent-NN, UQ-NN, and Phys-NN having higher χ^2 values than MaxEnt, they yield more accurate response functions, as evidenced by the entropy values. It is important to note that Ent-NN provides a smaller χ^2 value than Phys-NN, while UQ-NN is slightly less accurate. This behavior can be attributed mainly to the training process of UQ-NN, which involves introducing noise into the Euclidean responses, as discussed in Sec. III. As expected, introducing noise into the model inherently degrades its accuracy. On the other hand, it enhances the model's robustness, allowing it to propagate the uncertainties of the Euclidean responses in the reconstructed response functions.

Figure 3 displays the box plot of the S and χ^2 distributions for the one-peak (left column) and two-peak (right column) datasets obtained within the Ent-NN, UQ-NN, Phys-NN, and MaxEnt methods. Consistent with the results listed in Table I, the one-peak χ^2 and S distributions are narrower and

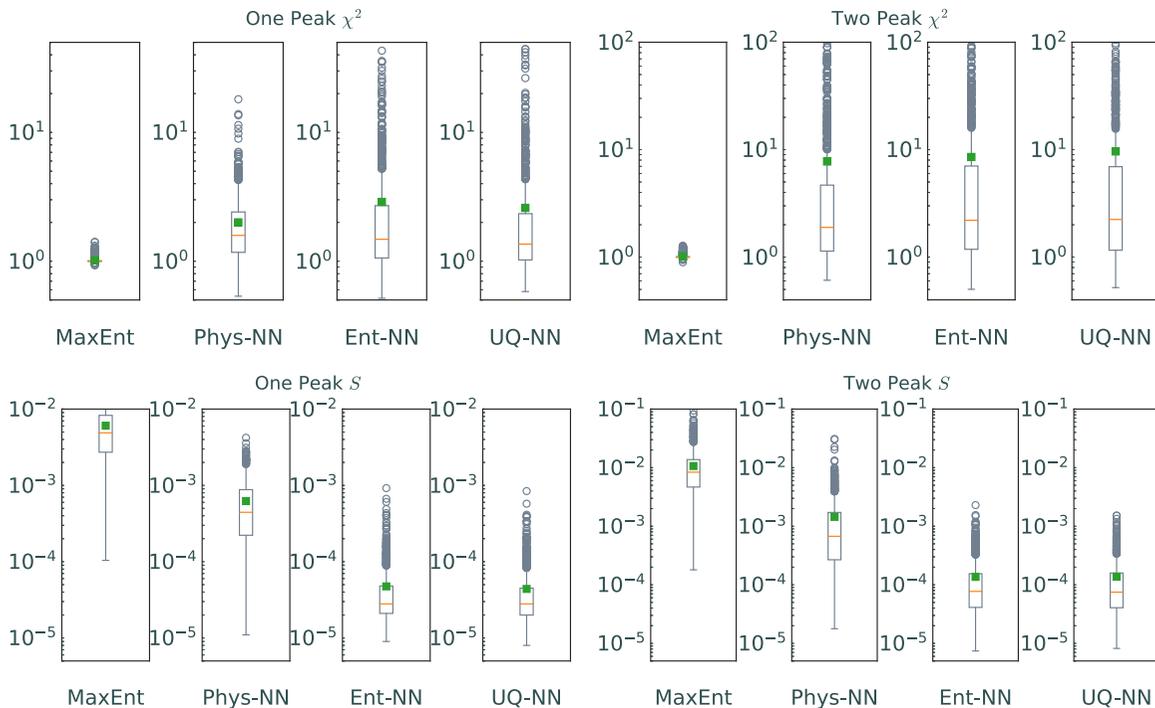


FIG. 3. Box plots of (top row) χ^2 and (bottom row) S for the one peak dataset (left) and two peak dataset (right) as obtained with Ent-NN, UQ-NN, Phys-NN, and MaxEnt. The line in the middle of the box denotes the median, and the box represents the range between the 25% and 75% quantiles. Whiskers cover the area between the 1% and 99% quantiles; data beyond these whiskers are outliers and are indicated by circles. The green squares represent the mean values of the scores.

centered on smaller values than the two-peak ones. Since Ent-NN, UQ-NN, and Phys-NN are trained to keep the reconstructed response function close to the original ones, we observe a much smaller spread of \bar{S} values compared with MaxEnt.

As previously discussed, historic MaxEnt naturally produces χ^2 values tightly clustered around 1. In contrast, the spread associated with Ent-NN and UQ-NN is more extensive, as evident in the scatter plots of Fig. 4. Correlations between the χ^2 values and S are noticeable in both neural network architectures, particularly for the two-peak dataset. This correlation provides a crucial advantage over MaxEnt, serving as a tool to assess the accuracy of the Laplace transform inversion. For instance, the outliers in χ^2 visible in the top right corners of all panels in Fig. 4 serve as clear indicators of imperfect reconstructions of the response functions.

Figure 5 provides a comparative illustration of the various inversion approaches. From the combined dataset, we select the *best* (left panels), *average* (central panels), and *worst* (right panels) reconstructed response functions based on their S values obtained within Ent-NN. Remarkably, not only the “best” and the “average” response functions but also the “worst” response functions reconstructed with Ent-NN exhibit closer agreement with the original ones compared to those obtained with historic MaxEnt. We also note that Ent-NN performs similarly to Phys-NN even in the worst case of Ent-NN, which does not necessarily correspond to the worst-case for Phys-NN. Notably, the uncertainty band of Ent-NN includes the original response even in this scenario.

Additionally, the Laplace transforms of the Ent-NN and UQ-NN response functions show excellent alignment with the original Euclidean responses: the χ^2 values are 0.37, 8.15,

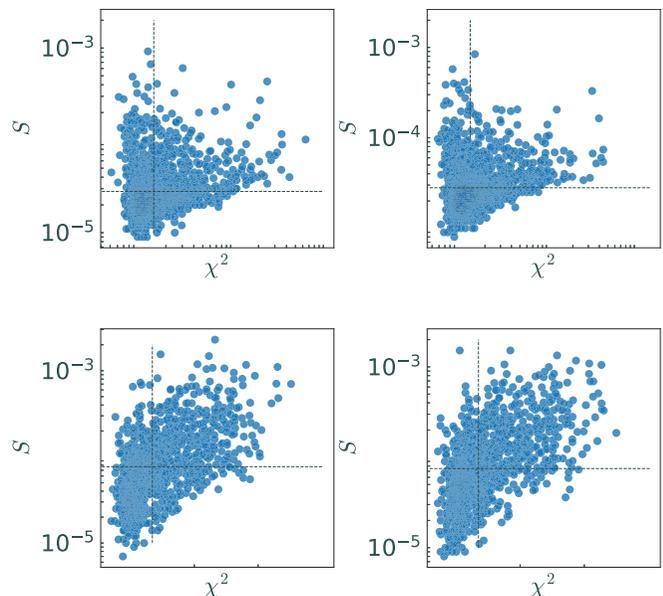


FIG. 4. Scatter plots of χ_E^2 versus S for Ent-NN (left column) and UQ-NN (right column). The top row corresponds to the one-peak dataset, while the bottom row shows results for the two-peaks dataset. Dashed lines indicate the median values $\bar{\chi}_E^2$ and \bar{S} .

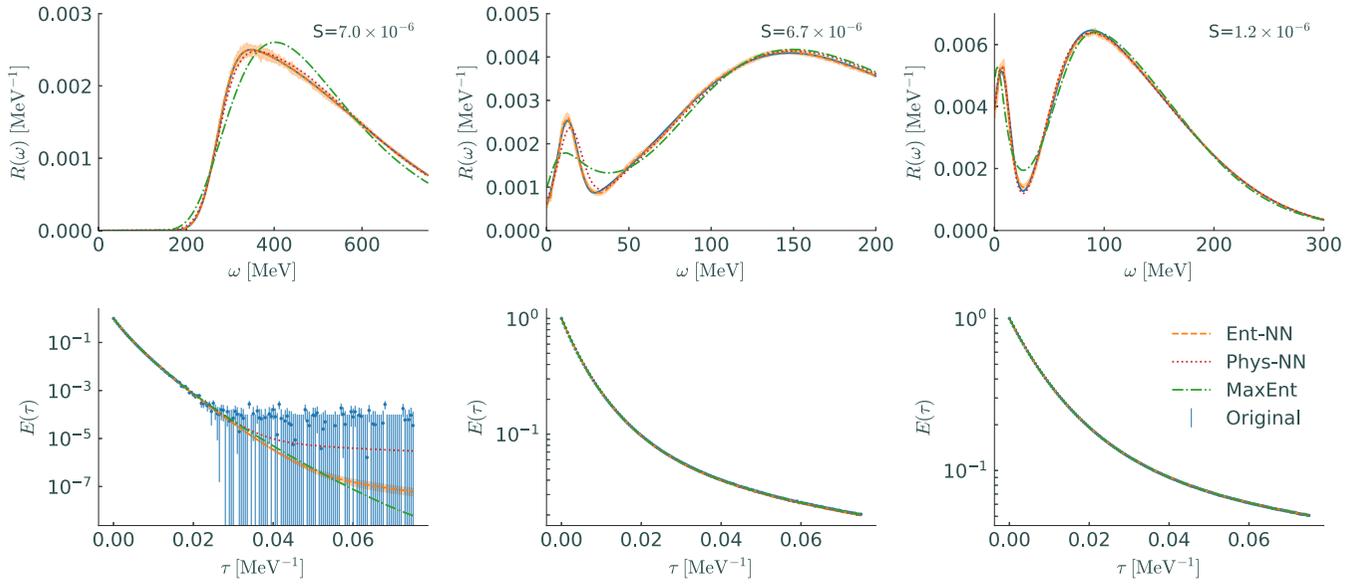


FIG. 5. Comparison between Ent-NN, Phys-NN, and MaxEnt for the combined dataset, from which we select the *best* (left panels), *average* (central panels), and *worst* (right panels) responses based on their S values obtained within Ent-NN. The top row displays the response functions and the bottom row the corresponding Euclidean responses. The shaded area around the Ent-NN results displays their epistemic uncertainty, estimated by training ten independently initialized neural networks.

and 342.27 for the best, average, and worst reconstructions, respectively.

B. Uncertainty quantification

As discussed in Sec. III, to perform uncertainty quantification, we generated 1000 Euclidean responses according to Eq. (18) for both the one- and two-peaks datasets. The corresponding reconstructed responses, distributed according to $P(\mathbf{R})$ of Eq. (16), are obtained applying UQ-NN to each of these Euclidean responses. Both the input noisy Euclidean

and reconstructed responses are displayed in Figs. 6 and 7 for the one- and two-peaks datasets, respectively, with varying degrees of noise in the input. The latter ranges from $\sigma = 10^{-4}$, corresponding to GFMC calculations of ^{12}C , to $\sigma = 10^{-2}$, which is appropriate for ongoing AFDMC calculations of ^{16}O Euclidean responses. Note that in the inference phase, as opposed to the training, we do not input the original Euclidean to UQ-NN, but just the noisy one, which is the only one available in real-world scenarios.

The key feature of our UQ-NN model is that it is able to capture the uncertainties present in the input data.

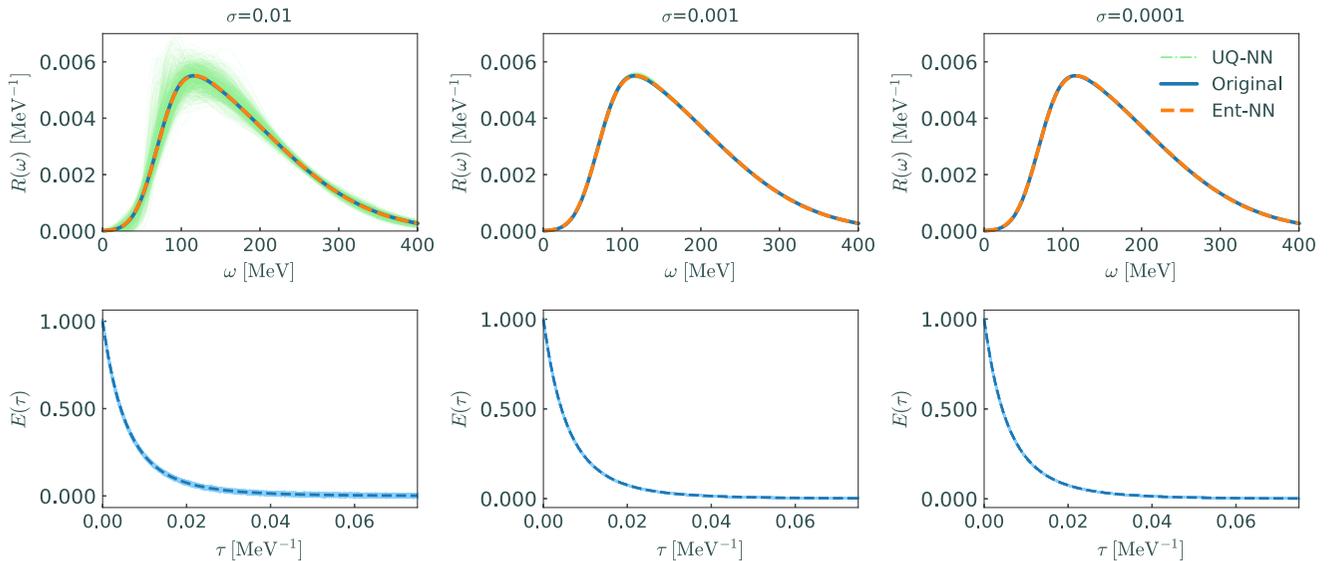


FIG. 6. Top row: one-peak responses obtained via the UQ-NN architecture (green band), compared to Ent-NN (dashed orange line) and the original response (blue solid line). Bottom row: input Euclidean responses with varying noise level.

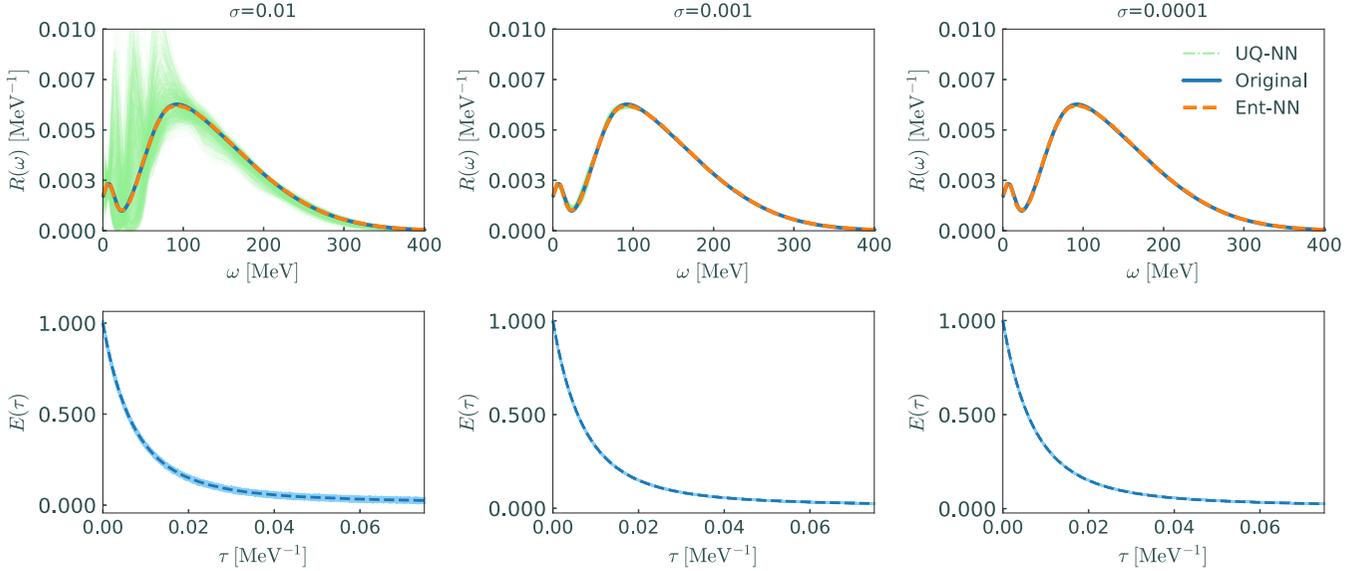


FIG. 7. Same as Fig. 6 for the two-peaks dataset.

In fact, the uncertainty in the reconstructed responses is generally proportional to the amount of statistical noise in the input Euclidean. This is already apparent in the one-peak dataset results; the green band around the reconstructed responses becomes larger as the noise level in the input increases. Notably, the original response function remains always enveloped within these bands, corroborating the accuracy of UQ-NN in both reconstructing the response function and in propagating the uncertainties of the input Euclidean.

Similar observations can be made for the two-peaks dataset. Here, however, we can observe some additional features of the UQ-NN responses. First, at the largest noise level $\sigma = 10^{-2}$, a three-peaks structure seems to emerge in the low- ω region, despite no response functions in the training dataset have more than two peaks. We ascribe the origin of this rich structure to the noise added to the Euclidean, which may yield a three-peak structure in some of the reconstructed $\hat{\mathbf{R}}(\theta)$. Secondly, for the lowest noise level, $\sigma = 10^{-4}$, UQ-NN fails to precisely capture the ω dependence of the original response function, even in the QE peak region. We checked that Ent-NN (and even Phys-NN) suffers from similar limitations. One possible reason for this behavior are numerical errors associated with numerically computing the Laplace transform—see Eq. (3)—when generating the training data set. The latter could be larger than the estimated 10^{-5} value, especially for responses with two peaks. Another possibility is the uncertainty inherent to the neural-network model, which includes the set of optimal parameters found in the training procedure and the training itself. To better estimate the latter, we plan on using a deep Bayesian neural network [45,46], which in the context of nuclear physics has proved reliable in predicting masses and radii of several nuclei across the nuclear chart, with quantified uncertainties.

V. CONCLUSIONS

Microscopic calculations of nuclear electromagnetic response functions are essential for connecting aspects of short- and long-range nuclear dynamics observed in electron-scattering experiments [13,47–51] to the underlying nuclear interactions and currents. Additionally, the availability of accurate electroweak (neutral and charge-current) response functions with quantified theoretical uncertainties is crucial for the success of the accelerator neutrino program [16–18,42], as nuclear cross section uncertainties are among the primary sources of systematic errors.

Over the past decade, the GFMC method has been extensively employed to compute electroweak response functions of nuclei with up to $A = 12$ nucleons, including one and two-body current operators consistent with the Hamiltonian generating correlations in the initial and final state of the reactions [19,20,52]. More recently, coupled cluster theory has achieved remarkable success in modeling longitudinal and transverse electromagnetic responses of nuclei as large as ^{40}Ca [21,53,54], retaining one-body current contributions only. In contrast to methods relying on harmonic-oscillator expansions [21], GFMC faces no challenges in handling high-resolution (or high-momentum) nuclear forces. These capabilities are indispensable for modeling the final state of reactions with momentum transfers above $q \simeq 400$ MeV. However, reconstructing electroweak response functions from GFMC imaginary-time propagators involves solving the notoriously ill-posed problem of inverting the Laplace transform.

In this study, we introduced two artificial neural network architectures suitable for approximating the inversion of the Laplace transform: Ent-NN and UQ-NN. A significant advantage over existing architectures, such as Phys-NN [32], is that both Ent-NN and UQ-NN utilize basis functions

tailored to the Laplace kernel, determined through its singular value decomposition [24]. We demonstrated their effectiveness by benchmarking Ent-NN and UQ-NN against Phys-NN and MaxEnt, using a substantial dataset comprising synthetic yet realistic data characterized by a broad quasielastic peak and a sharper elastic peak at lower energies. Ent-NN outperforms both Phys-NN and MaxEnt in terms of both metrics we considered. As a significant advance with respect to existing approaches, UQ-NN is designed to propagate the statistical errors of the Euclidean response through the response functions, which is critical for carrying out quantitative comparisons with experimental data.

The results presented in this work are particularly relevant for extending quantum Monte Carlo calculations to nuclei larger than ^{12}C , specifically to ongoing AFDMC calculations of electroweak response functions of ^{16}O . Standard MaxEnt suffers from severe limitations in this case for two main reasons. First, AFDMC suffers from significantly larger statistical noise than GFMC, primarily due to a stronger fermion sign problem [44]. Second, the electromagnetic responses of ^{16}O exhibit a rich low-energy structure, especially in the longitudinal channel, including elastic transitions and collective modes [55].

Notably, while our architectures were explicitly developed to approximate the inverse of the Laplace transform, they can be readily extended to different kernels, including the Lorentz

one. Consequently, they serve as valuable benchmarks for other inversion techniques, such as those based on expanding response functions on regularized ansatz [56], which require imposing the breakup threshold or employing appropriate expansions in Chebyshev polynomials [22].

ACKNOWLEDGMENTS

We express our sincere gratitude to R. B. Wiringa and C. Adams for their careful reading of the manuscript and their valuable suggestions. We are deeply grateful for the enlightening and stimulating discussions with P. Balaprakash, N. Rocco, and S. Wild, whose valuable insights significantly contributed to the development of this work. Additionally, we express our sincere thanks to G. Carleo for his unwavering support for this research direction. This work was supported in part by the U.S. Department of Energy (DOE), Office of Science, Offices of Advanced Scientific Computing Research and Nuclear Physics, by the Argonne LDRD program, and by the NUCLEI, FASTMath, and RAPIDS SciDAC projects under Contract No. DE-AC02-06CH11357. A.L. is also supported by DOE Early Career Research Program awards. We are grateful for the computing resources from the Joint Laboratory for System Evaluation and Leadership Computing Facility at Argonne.

-
- [1] B. R. Barrett, P. Navratil, and J. P. Vary, *Prog. Part. Nucl. Phys.* **69**, 131 (2013).
 - [2] G. Hagen, T. Papenbrock, M. Hjorth-Jensen, and D. J. Dean, *Rep. Prog. Phys.* **77**, 096302 (2014).
 - [3] H. Hergert, S. K. Bogner, T. D. Morris, A. Schwenk, and K. Tsukiyama, *Phys. Rep.* **621**, 165 (2016).
 - [4] A. Carbone, A. Cipollone, C. Barbieri, A. Rios, and A. Polls, *Phys. Rev. C* **88**, 054326 (2013).
 - [5] E. Epelbaum, H. Krebs, D. Lee, and Ulf-G. Meissner, *Phys. Rev. Lett.* **106**, 192501 (2011).
 - [6] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. Wiringa, *Rev. Mod. Phys.* **87**, 1067 (2015).
 - [7] B. Hu *et al.*, *Nat. Phys.* **18**, 1196 (2022).
 - [8] P. Gysbers *et al.*, *Nat. Phys.* **15**, 428 (2019).
 - [9] A. Roggero and J. Carlson, *Phys. Rev. C* **100**, 034610 (2019).
 - [10] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
 - [11] M. Schmitt and M. Heyl, *Phys. Rev. Lett.* **125**, 100503 (2020).
 - [12] A. Miessen, P. J. Ollitrault, F. Tacchino, and I. Tavernelli, *Nat. Comput. Sci.* **3**, 25 (2023).
 - [13] O. Hen, G. A. Miller, E. Piasetzky, and L. B. Weinstein, *Rev. Mod. Phys.* **89**, 045002 (2017).
 - [14] E. P. Segarra, A. Schmidt, T. Kutz, D. W. Higinbotham, E. Piasetzky, M. Strikman, L. B. Weinstein, and O. Hen, *Phys. Rev. Lett.* **124**, 092002 (2020).
 - [15] E. P. Segarra, J. R. Pybus, F. Hauenstein, D. W. Higinbotham, G. A. Miller, E. Piasetzky, A. Schmidt, M. Strikman, L. B. Weinstein, and O. Hen, *Phys. Rev. Res.* **3**, 023240 (2021).
 - [16] O. Benhar, P. Huber, C. Mariani, and D. Meloni, *Phys. Rep.* **700**, 1 (2017).
 - [17] T. Katori and M. Martini, *J. Phys. G* **45**, 013001 (2018).
 - [18] L. Alvarez-Ruso *et al.* (NuSTEC Collaboration), *Prog. Part. Nucl. Phys.* **100**, 1 (2018).
 - [19] A. Lovato, S. Gandolfi, J. Carlson, S. C. Pieper, and R. Schiavilla, *Phys. Rev. Lett.* **117**, 082501 (2016).
 - [20] A. Lovato, J. Carlson, S. Gandolfi, N. Rocco, and R. Schiavilla, *Phys. Rev. X* **10**, 031068 (2020).
 - [21] J. E. Sobczyk, B. Acharya, S. Bacca, and G. Hagen, *Phys. Rev. Lett.* **127**, 072501 (2021).
 - [22] J. E. Sobczyk and A. Roggero, *Phys. Rev. E* **105**, 055310 (2022).
 - [23] J. Carlson, J. Jourdan, R. Schiavilla, and I. Sick, *Phys. Rev. C* **65**, 024002 (2002).
 - [24] R. Bryan, *Eur. Biophys. J.* **18**, 165 (1990).
 - [25] M. Jarrell and J. Gubernatis, *Phys. Rep.* **269**, 133 (1996).
 - [26] A. Lovato, N. Rocco, and R. Schiavilla, *Phys. Rev. C* **100**, 035502 (2019).
 - [27] M. T. McCann, K. H. Jin, and M. Unser, *IEEE Signal Process. Mag.* **34**, 85 (2017).
 - [28] L.-F. Arsenault, R. Neuberger, L. A. Hannah, and A. J. Millis, *Inverse Probl.* **33**, 115007 (2017).
 - [29] H. Yoon, J.-H. Sim, and M. J. Han, *Phys. Rev. B* **98**, 245101 (2018).
 - [30] X. Xie, F. Bao, T. Maier, and C. Webster, *Discrete and Continuous Dynamical Systems* **15**, 877 (2022).
 - [31] R. Fournier, L. Wang, O. V. Yazyev, and Q. S. Wu, *Phys. Rev. Lett.* **124**, 056401 (2020).
 - [32] K. Raghavan, P. Balaprakash, A. Lovato, N. Rocco, and S. M. Wild, *Phys. Rev. C* **103**, 035502 (2021).

- [33] G. Shen, L. E. Marcucci, J. Carlson, S. Gandolfi, and R. Schiavilla, *Phys. Rev. C* **86**, 035503 (2012).
- [34] J. Golak, R. Skibiński, K. Topolnicki, H. Witała, A. Grassi, H. Kamada, and L. E. Marcucci, *Phys. Rev. C* **98**, 015501 (2018).
- [35] J. Carlson and R. Schiavilla, *Phys. Rev. Lett.* **68**, 3682 (1992).
- [36] S. Kullback and R. A. Leibler, *Ann. Math. Stat.* **22**, 79 (1951).
- [37] S. F. Gull and G. J. Daniell, *Nature (London)* **272**, 686 (1978).
- [38] J. Skilling, Classic maximum entropy, in *Maximum Entropy and Bayesian Method*, edited by J. Skilling (Springer Netherlands, Dordrecht, 1989), pp. 45–52.
- [39] W. Von Der Linden, R. Preuss, and V. Dose, in *Maximum Entropy and Bayesian Methods*, edited by W. von der Linden, V. Dose, R. Fischer, and R. Preuss (Springer Netherlands, Dordrecht, 1999), pp. 319–326.
- [40] M. Hohenadler, D. Neuber, W. von der Linden, G. Wellein, J. Loos, and H. Fehske, *Phys. Rev. B* **71**, 245111 (2005).
- [41] D. P. Kingma and J. Ba, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA* (ICLR, 2015).
- [42] L. A. Ruso *et al.*, [arXiv:2203.09030](https://arxiv.org/abs/2203.09030).
- [43] J. Zhang, *WIREs Comput. Stat.* **13**, e1539 (2021).
- [44] K. Schmidt and S. Fantoni, *Phys. Lett. B* **446**, 99 (1999).
- [45] R. Utama, J. Piekarewicz, and H. B. Prosper, *Phys. Rev. C* **93**, 014311 (2016).
- [46] Z. M. Niu and H. Z. Liang, *Phys. Lett. B* **778**, 48 (2018).
- [47] T. W. Donnelly and J. D. Walecka, *Annu. Rev. Nucl. Sci.* **25**, 329 (1975).
- [48] O. Benhar, D. Day, and I. Sick, *Rev. Mod. Phys.* **80**, 189 (2008).
- [49] S. Kegel *et al.*, *Phys. Rev. Lett.* **130**, 152502 (2023).
- [50] D. Adhikari *et al.* (PREX Collaboration), *Phys. Rev. Lett.* **126**, 172502 (2021).
- [51] D. Adhikari *et al.* (CREX Collaboration), *Phys. Rev. Lett.* **129**, 042501 (2022).
- [52] L. Andreoli, J. Carlson, A. Lovato, S. Pastore, N. Rocco, and R. B. Wiringa, *Phys. Rev. C* **105**, 014002 (2022).
- [53] J. E. Sobczyk, B. Acharya, S. Bacca, and G. Hagen, *Phys. Rev. C* **102**, 064312 (2020).
- [54] J. E. Sobczyk, B. Acharya, S. Bacca, and G. Hagen, *Phys. Rev. C* **109**, 025502 (2024).
- [55] S. Bacca, N. Barnea, G. Hagen, M. Miorelli, G. Orlandini, and T. Papenbrock, *Phys. Rev. C* **90**, 064619 (2014).
- [56] V. D. Efros, W. Leidemann, and G. Orlandini, *Few-Body Syst.* **26**, 251 (1999).