

Precise neural network predictions of energies and radii from the no-core shell model

Tobias Wolfgruber,^{1,*} Marco Knöll,¹ and Robert Roth^{1,2,†}

¹*Institut für Kernphysik, Fachbereich Physik, Technische Universität Darmstadt, Schlossgartenstr. 2, 64289 Darmstadt, Germany*

²*Helmholtz Forschungsakademie Hessen für FAIR, GSI Helmholtzzentrum, 64289 Darmstadt, Germany*

(Dated: July 1, 2024)

For light nuclei, *ab initio* many-body methods such as the no-core shell model are the tools of choice for predictive, high-precision nuclear structure calculations. The applicability and the level of precision of these methods, however, is limited by the model-space truncation that has to be employed to make such computations feasible. We present a universal framework based on artificial neural networks to predict the value of observables for an infinite model-space size based on finite-size no-core shell model data. Expanding upon our previous ansatz of training the neural networks to recognize the observable-specific convergence pattern with data from few-body nuclei, we improve the results obtained for ground-state energies and show a way to handle excitation energies within this framework. Furthermore, we extend the framework to the prediction of converged root-mean-square radii, which are more difficult due to the much less constrained convergence behavior. For all observables robust and statistically significant uncertainties are extracted via the sampling over a large number of network realizations and evaluation data samples.

I. INTRODUCTION

The challenge of solving the nuclear many-body problem with realistic Hamiltonians is at the core of nuclear theory. Throughout the past decades, the access to growing high performance computing (HPC) resources has revolutionized nuclear structure physics and caused a paradigm shift from computationally trivial to more demanding numerical models, giving rise to a range of successful *ab initio* methods such as the no-core shell model (NCSM) [1–4], the coupled-cluster (CC) theory [5], self-consistent Green’s function (SCGF) approaches [6], the in-medium similarity renormalization group (IM-SRG) [7], or quantum Monte Carlo (QMC) methods [8]. While they have proven to provide realistic results for a range of observables, their precision is limited by computational constraints.

In a concurrent development, modern machine learning algorithms have excelled in a variety of computational tasks such as pattern recognition, interpolation, or optimization and have naturally sparked the interest of researchers in many fields including nuclear structure physics. This has led to various applications [9, 10] either through exploiting machine learning features as done in neural-network quantum state techniques [11–14], purely data driven approaches [15, 16], or as supplemental tools for the aforementioned many-body methods [17–19]. Focusing on the latter, machine learning is of particular interest when it comes to extending the reach of *ab initio* methods by circumventing their computational limitations. As most of them are based on a basis expansion of the many-body problem, a major restraint are the rapidly growing model spaces that quickly exceed the capabilities of even the largest HPC clusters available. Due to the nature of *ab initio* methods, which employ systematically improvable truncations, the errors induced by these truncations can, in principle, be controlled. However, accurately quantifying

them or extrapolating results from finite model spaces to the full many-body Hilbert space are extremely challenging tasks. While ground-state energies are accessible with conventional extrapolation schemes based on exponential fits [3, 20, 21], there are no widely established methods for other observables such as radii.

As shown in recent applications to NCSM or CC calculations of ground-state observables, artificial neural networks (ANNs) are a promising tool to tackle this challenge [17–19]. The convergence behavior of these methods w.r.t the model space size is controlled by a model-space truncation parameter and the harmonic-oscillator (HO) frequency of the underlying single-particle basis.

The ANN models for extrapolating such data separate into two conceptually different categories. The first category attempts to emulate the functional dependence of the observable on the model-space size or, more precisely, the truncation parameter and the HO frequency, thus, directly replacing exponential or polynomial extrapolation methods with an ANN [17, 18]. Approximations for the full Hilbert space are then given by the networks prediction for very large values of the model-space truncation parameter. While this provides robust extrapolations for well balanced sets of training data, the applicability of such an ANN is, by construction, limited to the nucleus, interaction, state, and observable it was trained for. Thus, the computation of a large and balanced sets of training data as well as the ANN training process itself has to be repeated for each nucleus, interaction, state, and observable under consideration. This is a major bottleneck as the generation of training data becomes very costly. Moreover, the ANNs are used for a true extrapolation, i.e., an evaluation in an input parameter regime they were never trained for.

We have introduced a different approach in Ref. [19], which is more akin to a pattern recognition task. In our ANN model, the observable in the full Hilbert space is directly predicted from a set of converging sequences in small model spaces, which share the same limit. Hence, the ANN attempts to capture the convergence pattern and predict the converged value, similar to how an experienced practitioner would estimate it. Since we need to specify the fully converged values for the

* tobias.wolfgruber@physik.tu-darmstadt.de

† robert.roth@physik.tu-darmstadt.de

training process, we can only use NCSM data from very light systems with $A \leq 4$ for the training set. Despite this limitation, we can construct an extensive library of training data for a specific observable by varying the underlying Hamiltonian in addition to the nucleus and model-space parameters. In this way the training data encompasses very diverse convergence patterns so that, with a proper normalization, the ANN has learned all possible convergence patterns during the training. This makes this ANN universally applicable to different nuclei, interactions and states. This universality is a great advantage. The generation of training data and the training only happens once and the resulting ANNs are then simply applied with small sets of NCSM evaluation data for the systems of interest.

While our previous work [19] has demonstrated the capabilities of this universal ANN model for predicting ground-state energies, this paper focuses on improvements of the precision of the ANN prediction and the generalization to other observables, particularly excitation energies and root-mean-square (rms) radii. The description of these observables is more challenging because they show a much larger variety of non-monotonous convergence patterns, since they are not protected by the variational principle. We will introduce a simple normalization scheme that improves the precision of predictions of converged ground-state and excited-state energies and enables the robust prediction of converged radii. For all observables we will extract statistically meaningful uncertainties for the ANN predictions. We will demonstrate the application of the ANNs to a range of p-shell nuclei up to ${}^9\text{Be}$.

II. NO-CORE SHELL MODEL

While these ANN approaches can be adapted to various many-body methods, we focus on the NCSM, which solves the many-body stationary Schrödinger equation as a matrix eigenvalue problem

$$\sum_j \langle \phi_i | H | \phi_j \rangle \langle \phi_j | \psi_n \rangle = E_n \langle \phi_i | \psi_n \rangle \quad \forall i, \quad (1)$$

by expanding the Hamiltonian H as well as the eigenstates $|\psi_n\rangle$ in a complete basis. The basis $\{|\phi_i\rangle\}$ is a set of Slater determinants of single-particle basis states, in our case eigenstates of the HO. To numerically solve this matrix eigenvalue problem, a truncation to a finite model space is employed. The so-called N_{max} truncation limits the total number of HO excitation quanta. The convergence behavior as function of N_{max} depends on the chosen HO frequency $\hbar\Omega$ of the single-particle basis, though in the limit $N_{\text{max}} \rightarrow \infty$ this dependency vanishes and the exact solution is recovered.

Observables other than the energies E_n are easily accessible in the NCSM by constructing the matrix representation of the corresponding operator in the many-body basis and evaluating the expectation value with the calculated eigenvectors. Since the NCSM is a variational method, the energy eigenvalues converge monotonously from above to the exact value with increasing N_{max} . This powerful constraint on the convergence pattern does, however, not hold for other observables,

e.g. radii, leading to more complex and less constrained convergence patterns.

The Hamiltonian employed in the NCSM calculations consists of a kinetic energy contribution and inter-nucleon interactions, typically derived from chiral effective field theory (EFT). There are several families of chiral nucleon-nucleon (NN) plus three-nucleon (3N) interactions with different order in the chiral power-counting and different cutoffs available nowadays, which exhibit different NCSM convergence behaviors. Even for the softest, i.e., fastest converging, of the interactions very large N_{max} are required to reach convergence. Therefore, an additional pre-processing of the Hamiltonian through a similarity renormalization group (SRG) transformation is used to accelerate the model space convergence [20, 22–25]. The SRG transformation introduces the continuous flow parameter α which provide direct control over the convergence rate.

For few-body systems, which include ${}^2\text{H}$, ${}^3\text{H}$, and ${}^4\text{He}$, we can use the Jacobi-NCSM [26], which is a efficient formulation of the NCSM for very light nuclei. It uses a relative basis instead of the typical m -scheme, which allows us to obtain fully converged calculations at minimal computational expense, ideal for generating large sets of training data. For heavier nuclei in the p-shell, as they are used as test cases for the network evaluation, we use the standard m -scheme NCSM.

III. ARTIFICIAL NEURAL NETWORKS

For our purpose, artificial neural networks are an ideal tool, as they excel at capturing patterns, especially in cases where classical algorithms are hard to implement or where the exact correlations between data points are unknown. At their core, ANNs roughly emulate biological neural networks as present in the brain. The type of network we use is called a dense feed-forward neural network. It consists of neurons with inputs and outputs, which are organized in layers. In this context dense means that each neuron's input is connected to the output of every neuron in the preceding layer. Based on these inputs, each neuron then calculates an output, which is fed to every neuron in the subsequent layer. For example, the i -th neuron in the n -th layer computes its output $x_n^{(i)}$ according to

$$x_i^{(n)} = \sigma \left(\sum_{j=0}^{N_{n-1}} x_j^{(n-1)} w_{ji}^{(n)} + b_j^{(n)} \right), \quad (2)$$

where N_{n-1} is the number of connected neurons from the previous layer, $w_{ji}^{(n)}$ and $b_j^{(n)}$ are the weights and biases which are optimized during training, and σ is the activation function which rescales the output and in general introduces nonlinearities. The weights and biases are initialized randomly and are optimized during the training process via supervised learning. This means that every network input in the training data is labeled with the desired output.

The training process starts with a feed-forward pass, where a batch of training samples are evaluated and a combined deviation of the network's predictions from the desired results is

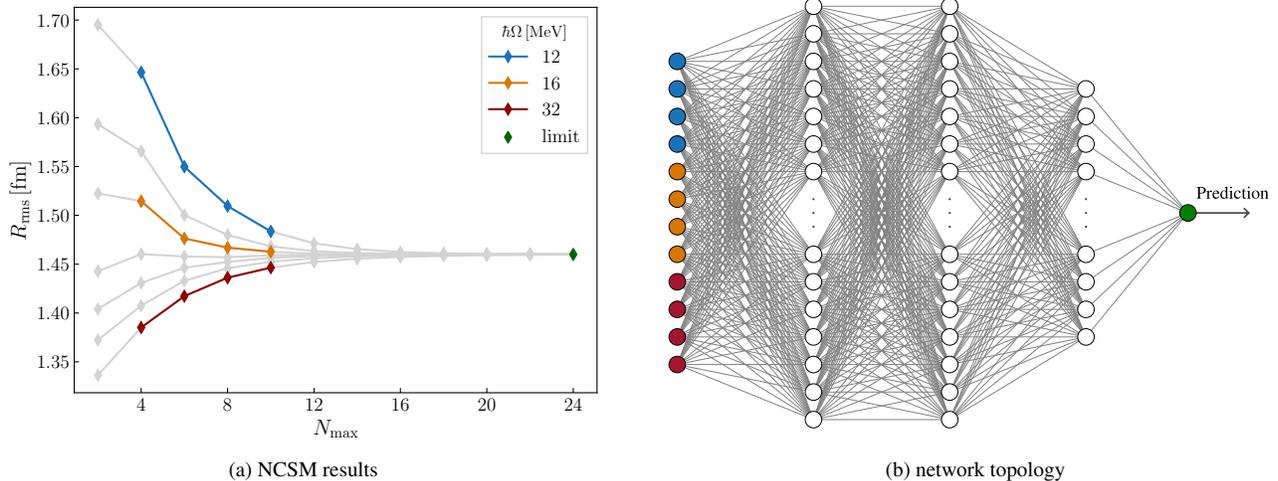


FIG. 1. (a) Exemplary Jacobi-NCSM results for the mass rms radius of ${}^4\text{He}$. All training data sequences for this combination of nucleus and interaction are shown in gray. Highlighted in color is one possible training sample, consisting of four radius results for consecutive N_{max} values for 3 different HO frequencies, and the fully converged target value (green). (b) Schematic representation of the fully connected feed-forward network topology, where the colors of the input and output layer correspond to the colors of the sample in (a).

calculated with a so-called loss function. The following back propagation step optimizes the weights and biases based on the gradient of the loss function. Repeating this process leads to an iterative improvement of the parameters of the network. A validation data set that is distinct from the training data is used to probe the performance of the network and the training progress and is usually used to adjust the step size in the gradient descent, commonly called learning rate. In the end, one should obtain a network that captures the patterns in the training data, leading to accurate predictions on the training data as well as previously unseen evaluation data.

So far we did not specify the choice of the so-called hyperparameters like the activation function, loss function, or back propagation algorithm. These choices can have a significant impact on the performance on both the training efficiency as well as on the quality of the networks themselves. We achieved the best results with a rectified linear unit (ReLU) [27] activation function, a mean-square error (MSE) loss function, and the AdamW back-propagation algorithm [28] in combination with an adaptive learning rate scheduler, that halves the learning rate when the loss plateaus for two consecutive epochs, and a batch size of 512. Note that, while the overall choice of hyperparameters is crucial for the performance of the ANNs, the predictions are very robust under reasonably small changes.

IV. CONCEPTS AND NETWORK DESIGN

As already mentioned, artificial neural networks are ideally suited for situations where data is correlated, but where the patterns are either too varied or the problem is too hard to write down a functional dependence. This is exactly the case for NCSM calculations, where the convergence with N_{max} clearly follows a certain pattern, that is hard to quantify. The

idea is to show the network multiple sequences of data points which all must converge to the same value and have the network output a prediction for the fully converged value of these sequences. An illustration of converging NCSM sequences for the rms-radius of ${}^4\text{He}$ is shown in Fig. 1a.

To construct such ANNs, we first need to train them on realistic NCSM convergence data. Since we use supervised learning, the exact, fully converged observable is a requirement for every sequence in the training data, which limits us to using few-body systems for the training as only these can be converged to the required precision. However, since the convergence patterns in heavier p-shell nuclei are very similar to those in few-body systems, ANNs trained on the latter can provide precise predictions for a broad range of p-shell nuclei. The main benefit of this method is that the training only needs to be done once and the trained ANNs can subsequently be applied, as a universal tool, to convergence sequences from any NCSM calculation.

Note that each observable needs its own set of networks, as the convergence pattern can vary significantly between different observables, though strongly related observables such as the ground-state energy and excited-state energies or rms radii and point-proton rms radii can use the same set of networks.

One important aspect we have not touched on yet is the exact design of the ANN including the structure of the data samples, which defines the input and output layers, and the topology, which determines the number of hidden layers. For our application, the input for the network consists of X sequences of NCSM calculations for different $\hbar\Omega$ which contain L values of a specific observable for subsequent values of N_{max} . This sets the input layer to be $X \cdot L$ neurons wide, while the output layer consists of a single neuron, which is supposed to be the prediction for the fully converged value of the observable. We found that three hidden layers with sizes $4X \cdot L$, $4X \cdot L$, and $8X$ work very well for our use case. Fewer or much smaller

hidden layers lead to less precise predictions, while adding more hidden layers or increasing their size yield diminishing returns, adding computational cost and increasing the risk of overfitting for only very marginal gains in precision. Figure 1 visualizes both a training sample (a) and the network topology (b) for the input dimensions of $X = 3$ and $L = 4$, which is the structure used for all results presented in this work. Note that the input layer only accepts a subset of the available data in the NCSM sequences—there are typically more than $X = 3$ different oscillator frequencies and more than $L = 4$ different N_{\max} values available. This opens up the opportunity for a random sampling of the input data, which will be important for the training and statistical evaluation discussed below.

In order to evaluate the performance of our networks, it is helpful to compare them with traditional extrapolation techniques. For energy extrapolations we compare to an exponential fit scheme as described in Ref. [25] for both the ground-state energy and excited-state energies. Radii are, due to being less constrained, more difficult to extrapolate. As there is no widely used algorithm for the extrapolation of radii, we forgo a comparison with a classical extrapolation technique for this observable. An interesting point for a future comparison are the effective-theory-based infrared extrapolations for the NCSM [29, 30]. The applicability of these methods relies on specific and often untypical combinations of N_{\max} and $\hbar\Omega$, which require a large number of dedicated NCSM calculations and are not directly possible with the evaluation data used here.

V. TRAINING DATA AND STATISTICAL EVALUATION

Our training data consists of ^2H , ^3H , and ^4He Jacobi-NCSM calculations up to $N_{\max} = 50, 40, \text{ and } 24$, respectively, for seven HO frequencies $\hbar\Omega = 12, 14, 16, 20, 24, 28, \text{ and } 32$ MeV. We use the non-local NN+3N interactions from chiral EFT at orders N^2LO , N^3LO , and $\text{N}^4\text{LO}'$ and for cutoffs $\Lambda = 450, 500, \text{ and } 550$ MeV introduced in [31, 32]. We applied these interactions bare and SRG evolved with three different flow parameters of $\alpha = 0.02, 0.04, \text{ and } 0.08$ fm⁴. All in all, we obtain 756 converging sequences. To reduce artifacts in the training set, we exclude $N_{\max} = 0$ data and remove sequences that are not sufficiently converged, which means that the value of the observable at the largest available N_{\max} shows a deviation of more than 5% from the fully converged value. The data that is left forms the basis from which we can construct samples for training. Each sample consists of a subset of the training data with a specific number of randomly selected HO frequencies and a specific window in N_{\max} , together with the fully converged value of the observable of choice.

We construct three disjoint subsets of samples: a large set of training samples, and two smaller sets of test and validation samples. The training consists of 20 epochs, iterating through all training samples with a batch size of 512 each time, monitoring the average loss via the test set at every cycle. This information is used to adaptively adjust the learning rate, which is initialized to 0.001. We evaluate the quality of the fully trained networks by measuring their performance on the vali-

dation set and discard poorly performing networks.

When applying the trained ANNs to unseen data we rely on a statistical evaluation for which we first construct evaluation samples analogously to the sample generation done for the training data, i.e., we construct all possible samples of $X = 3$ HO frequencies and their permutations with values for the observable at $L = 4$ consecutive N_{\max} . In addition to evaluating multiple samples, we also evaluate these samples with 1000 networks, that have been initialized differently and trained individually, in order to suppress any errors from a single ANN. This multitude of predictions is then incorporated into a histogram that describes the distribution of the individual predictions, which usually shows a strong, single-peak structure that is very similar to a normal distribution. For simplicity, we fit a Gaussian to the distribution to obtain the mean and standard deviation of the dominant peak structure. While the extracted mean can be understood as the final ANN prediction for the given observable, the width of the distribution provides an uncertainty measure that incorporates not only all information from the evaluation data, but also accounts for the uncertainty of the individual ANNs. Hence, we can provide precise predictions, which are robust against outliers, along with reliable 1σ uncertainties.

To be able to gauge the quality and consistency of the predictions, we group samples by the highest N_{\max} value they contain, denoted by \mathcal{N}_{\max} . Higher \mathcal{N}_{\max} correspond to larger model spaces and, therefore, better converged sequences, which should improve the ANN predictions of the fully converged value. To emulate the situation in heavier p-shell nuclei where we see the largest potential for this method, we artificially limit the information for the evaluation in this work to $N_{\max} = 12$ or less. This corresponds to evaluation samples with $\mathcal{N}_{\max} \leq 12$. With increasing \mathcal{N}_{\max} we expect the distribution of network predictions to decrease in width, corresponding to a smaller uncertainty, as well as improve in accuracy. Furthermore, the predictions should be compatible with each other within their uncertainty bands.

VI. INPUT MODES

There is some freedom when it comes to formatting the data that is fed into the input layer of the networks. We can formalize this by looking at the networks like a mapping $M : S \rightarrow O^\infty$, where O^∞ is the converged observable and S is the input sample. In our previous work [19], we compared two distinct input modes. In the ABS mode, the data for a given observable, e.g. the ground-state energy, is simply fed in as-is. With our chosen network topology, a single ABS sample takes the shape

$$S_{\text{ABS}}^{\mathcal{N}_{\max}} = (O_{\hbar\Omega_1}^{\mathcal{N}_{\max}-6}, O_{\hbar\Omega_1}^{\mathcal{N}_{\max}-4}, O_{\hbar\Omega_1}^{\mathcal{N}_{\max}-2}, O_{\hbar\Omega_1}^{\mathcal{N}_{\max}}, O_{\hbar\Omega_2}^{\mathcal{N}_{\max}-6}, O_{\hbar\Omega_2}^{\mathcal{N}_{\max}-4}, O_{\hbar\Omega_2}^{\mathcal{N}_{\max}-2}, O_{\hbar\Omega_2}^{\mathcal{N}_{\max}}, O_{\hbar\Omega_3}^{\mathcal{N}_{\max}-6}, O_{\hbar\Omega_3}^{\mathcal{N}_{\max}-4}, O_{\hbar\Omega_3}^{\mathcal{N}_{\max}-2}, O_{\hbar\Omega_3}^{\mathcal{N}_{\max}}).$$
 (3)

This is the naïve choice and gives the network all available information, but also the maximum amount of flexibility to

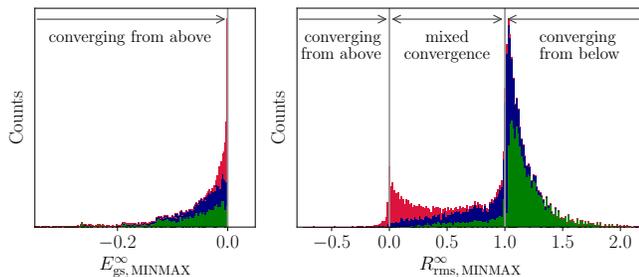


FIG. 2. Histogram of the dimensionless target values of all samples in the MINMAX input mode for ground-state energies (left) and mass rms radii (right), separated by color: ^2H (green), ^3H (blue), and ^4He (red). Values left of zero (right of one) belong to samples exclusively converging from above (below), while target values between zero and one belong to samples with a mixed convergence pattern.

look for patterns in the input data. A clear disadvantage is, however, that this can lead to undesired training results, as the networks just have to reproduce the few results for the nuclei in the training data without learning the convergence pattern. Previous work has further shown, that this input mode induces a dependency on the respective energy range, which distorts the predictions in nuclei beyond the training set. To circumvent this, we have introduced an additional step of scaling and shifting the data before feeding it to the network to increase variety in the target values and convergence patterns. This slightly alleviates the effects of the large gap between the target value of the heaviest training nucleus and the heavier nuclei in the evaluation set.

The other mode we introduced in Ref. [19] is called DIFF and feeds the differences between consecutive N_{max} steps into the networks. This inherently alleviates the difference in scales between the nuclei in the training and the evaluation set, though shifting the data obviously had no more effect so that one source for inflating the initial training data pool was eliminated. Both of these modes were explored with the ground-state energy as the target observable (for details see Ref. [19]). Unfortunately, both exhibit significant shortcomings when applied to an observable with a different convergence structure like the radius.

This leads us to look into another input mode called min-max normalization, which we refer to as MINMAX. The min-max normalization is a common normalization technique for machine learning applications. In general, it normalizes the input data per sample to an interval $[a, b]$, though we exclusively choose an interval of $[0, 1]$ for our network input. A single MINMAX sample can be defined via the ABS sample and takes the form

$$S_{\text{MINMAX}}^{N_{\text{max}}} = \frac{S_{\text{ABS}}^{N_{\text{max}}} - \min(S_{\text{ABS}}^{N_{\text{max}}})}{\max(S_{\text{ABS}}^{N_{\text{max}}}) - \min(S_{\text{ABS}}^{N_{\text{max}}})}. \quad (4)$$

To recover the value of the predicted observable O^∞ from the network output O_{MINMAX}^∞ , the inverse transformation has to be

applied:

$$O^\infty = O_{\text{MINMAX}}^\infty \left(\max(S_{\text{ABS}}^{N_{\text{max}}}) - \min(S_{\text{ABS}}^{N_{\text{max}}}) \right) + \min(S_{\text{ABS}}^{N_{\text{max}}}). \quad (5)$$

The goal is to get rid of scale dependencies and reduce the information to the network in such a way that the recognition and extrapolation of the convergence pattern is the only viable strategy for the network due to a lack of other clues or reference points contained within the data.

Figure 2 shows the distribution of the MINMAX-normalized target values of all training samples for ground-state energies and radii. The colors correspond to the three training nuclei. Both normalized target values are dimensionless, and the position on the x -axis determines if the value belongs to a sample that converges exclusively from above (below 0), exclusively from below (above 1), or one that has a mixed convergence (between 0 and 1). As ground-state energies are constrained by the variational principle, all samples must naturally converge from above. Furthermore, the ground-state energy usually converges rather fast, leading to most target values being grouped closely below zero. Radii on the other hand are, as previously mentioned, not very well constrained. The direction of convergence mainly depends on the specific HO frequency in conjunction with the nucleus and the N_{max} window. For both ^2H and ^3H , samples that converge from below dominate in our training set, while ^4He appears to be much more balanced in that regard. However, we expect this to be a good match for the situation of most nuclei we are interested in, as a typical energy-guided selection of HO frequencies leads to the majority of sequences for the rms-radius to converge from below. For the training, we generate 1 000 000 training samples for the ABS and DIFF input modes via random scaling and shifting, while we are limited to about 350 000 native samples for MINMAX.

The performance of the MINMAX input mode is compared to ABS and DIFF in Fig. 3 for the ^4He ground-state energy and mass rms-radius. The left-hand panels show the evaluation data for both observables, while the right-hand panels shows histograms of the networks' predictions together with the fitted Gaussian for all 3 input modes at increasing N_{max} . The fully converged value obtained from large N_{max} is depicted as a horizontal green line.

Starting with the ground-state energies, all three input modes perform well, though ABS gives slightly overbound predictions, while both DIFF and MINMAX are very accurate even at $N_{\text{max}} = 8$. DIFF gives very small uncertainties compared to both other modes, though the MINMAX uncertainties decrease significantly with increasing N_{max} down to approximately the level of DIFF, while the ABS uncertainties remain rather large. As the DIFF uncertainties can tend to be somewhat underestimated, MINMAX seems to be a very good alternative, retaining the accuracy of DIFF with the slightly more realistic uncertainties. The radius results in the lower panel show a much larger discrepancy between the different modes. Both ABS and DIFF give rather inaccurate results, especially for smaller N_{max} . The uncertainties of both modes also remain fairly large. The MINMAX input mode on the other hand is already extremely accurate at $N_{\text{max}} = 8$, and

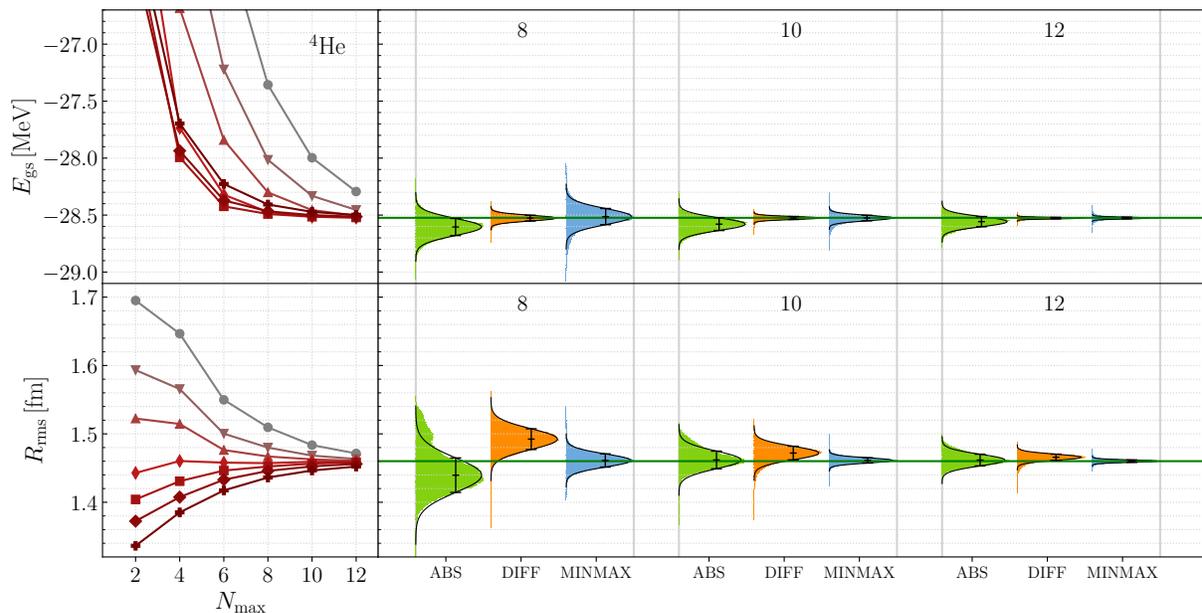


FIG. 3. Input data and network predictions for the ground-state energy and mass radius of ${}^4\text{He}$. The left column shows the input data, consisting of NCSM calculations for $\hbar\Omega = 12, 14, 16, 20, 24, 28,$ and 32 MeV (gray to red). The right column shows a histogram of the predictions of the networks using the ABS, DIFF, and MINMAX input modes as well as the fitted Gaussian and its mean and standard deviation. The three segments labeled 8, 10, and 12 correspond to the N_{max} of the evaluation samples.

subsequent N_{max} increase the precision in a realistic and systematic way when compared to the raw data on the left.

Though not shown here, the different input modes show a very similar behavior for other nuclei as well. This means that MINMAX is clearly superior for radii, and at least comparable if not better than both other input modes for ground-state energies. Therefore, we will exclusively use the MINMAX input mode for all following investigations.

VII. RESULTS FOR GROUND-STATE AND EXCITATION ENERGIES

In order to evaluate the performance of the networks, we will start with a look at ground-state energies of the few-body systems ${}^2\text{H}$, ${}^3\text{H}$, and ${}^4\text{He}$, which are also used in the train process. It is important to note that all evaluations shown in this work use an entirely different family of interactions than the non-local chiral interaction used during training. The evaluation data is obtained with a semi-local NN+3N interaction derived from chiral EFT at $N^2\text{LO}$ with a cutoff of $\Lambda = 450$ MeV [25, 33], SRG evolved with a flow parameter of $\alpha = 0.08$ fm 4 . This ensures maximum independence of training and evaluation data, even if the evaluation is performed for isotopes contained in the training data. The specific selection of HO frequencies for the NCSM evaluation data is given in the caption of the respective figure.

Figure 4 shows results for the few-body nuclei, the corresponding numerical values are summarized in Table I. As before, the left-hand panels show the raw evaluation data, while the right-hand panels show the distribution of the pre-

dictions together with a fitted Gaussian that defines the mean and standard deviation. The dotted lines indicate the variational minimum for $N_{\text{max}} = 8, 10,$ and 12 truncating the data fed into the networks. The results of traditional extrapolations at $N_{\text{max}} = 8, 10,$ and 12 are given in red for comparison. The green line is the fully converged result obtained at very large N_{max} .

As expected, the networks show excellent agreement with the fully converged results and are well below the variational minimum except when the sequences are already well converged. The prediction using the smallest $N_{\text{max}} = 8$ is in all 3 cases almost exactly on the fully converged value, indicating that the networks can recognize the convergence pattern very well and give accurate predictions. The uncertainties of the predictions get systematically smaller with increasing N_{max} , and all predictions are well within the preceding uncertainties indicating over-all consistency of the predictions and the uncertainty estimated.

Compared to the classical extrapolation, the networks are generally more accurate and precise in cases where the energy is not converged yet, which is the case for ${}^2\text{H}$ and ${}^3\text{H}$, especially at smaller N_{max} . Only the exceptionally well bound ${}^4\text{He}$ nucleus converges so fast, that there is no advantage compared to a classical extrapolation. It should be noted that for ${}^2\text{H}$ at $N_{\text{max}} = 10$ the uncertainty estimation of the classical extrapolation breaks down (specifically the exponential fit for $\hbar\Omega = 12$ MeV) due to a deuteron-specific artifact in the convergence sequence, resulting in an unreasonably large uncertainty.

Looking at slightly heavier nuclei, we are not only interested in the ground-state energy, but also in excited-state ener-

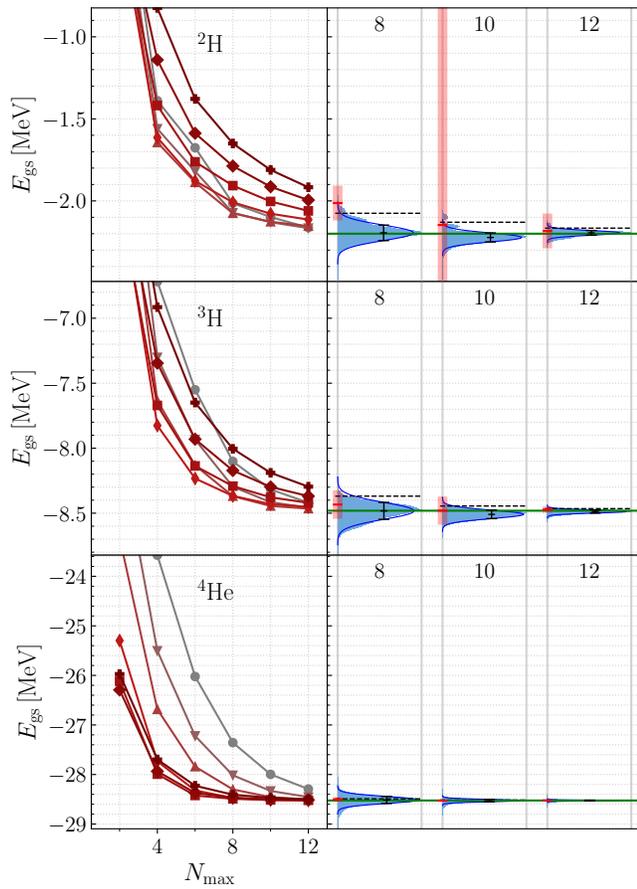


FIG. 4. Input data and network predictions of the ground-state energy for ${}^2\text{H}$, ${}^3\text{H}$, and ${}^4\text{He}$. The left hand panels show the NCSM input data for $\hbar\Omega = 12, 14, 16, 20, 24, 28,$ and 32 MeV (gray to red), while the right hand panels show a histogram of the network predictions together with a fitted Gaussian, which gives the mean and $1\text{-}\sigma$ uncertainty. The variational minimum at each N_{max} is given by the black dashed lines, the classical extrapolation results in red with the associated uncertainty, and the fully converges value is given by the green line.

gies or excitation energies. The energy of excited states can be predicted in exactly the same way using the same pre-trained networks as for the ground-state energy extrapolation, feeding sequences of the converging excited state to the networks as inputs instead. The upper panel of Fig. 5 shows the NCSM data for the 1^+ ground and 3^+ excited state of ${}^6\text{Li}$, with the corresponding network predictions on the right. There is no discernible difference in the performance of the networks for the two states, for both states the predictions are very stable and consistent with each other. This is possible since there is no qualitative difference in the convergence behavior of bound excited states compared to the ground state.

To give predictions for excitation energies with our networks, we evaluate both the ground and the excited-state energy separately, and subsequently subtract the predictions sample-wise, i.e., we compute the difference between the predicted ground-state energy and the predicted excited-state energy of for each sample consisting of $X = 3$ different HO fre-

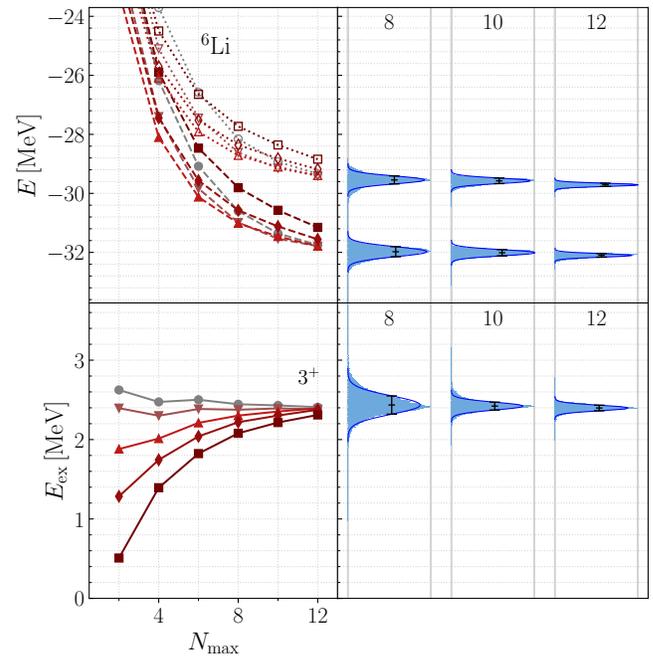


FIG. 5. Input data and network predictions for the 1^+ ground and 3^+ excited state of ${}^6\text{Li}$ (upper panel) and for the excitation energy of the 3^+ excited state (lower panel) with $\hbar\Omega = 14, 16, 20, 24,$ and 28 MeV (gray to red).

quencies and $L = 4$ subsequent model space truncations N_{max} . The resulting distribution tends to be approximately normal-shaped, so that we can once again simply fit a Gaussian to obtain the prediction for the excitation energy and its uncertainty from the mean and standard deviation the Gaussian.

The lower panels of Fig. 5 show the result of the above mentioned procedure for the 3^+ excitation energy on the right-hand side. The excitation energy sequences in the left-hand panel are only shown to give an impression of the convergence of the actual observable, as they are not used as input data for the networks. We observe that the excitation energy predictions are remarkably stable for the different N_{max} , given that this observable is a lot more challenging to predict, as slight variations in the prediction of the absolute energies has larger effects on the excitation energy. For this reason, we use the aforementioned sample-wise subtraction method instead of just taking the difference of the whole distributions. This way, we can make exploit of correlations in the convergence pattern of different states for the same HO frequency. Typically, this improves the prediction and leads to smaller uncertainties.

Figures 6 to 9 as well as Table I give a comprehensive overview of the performance of the networks for ground-state energies and the lowest few excitation energies of ${}^6\text{Li}$, ${}^7\text{Li}$, ${}^8\text{Li}$, and ${}^9\text{Be}$. These NCSM calculations for p-shell nuclei have been discussed in detail in Ref. [34] and we reuse these NCSM outputs here to benchmark the performance of the networks [35]. In these figures, the ground-state energy is shown in the lower panel, while the excitation energy results are in the upper panel. Additionally, classical extrapolations for both

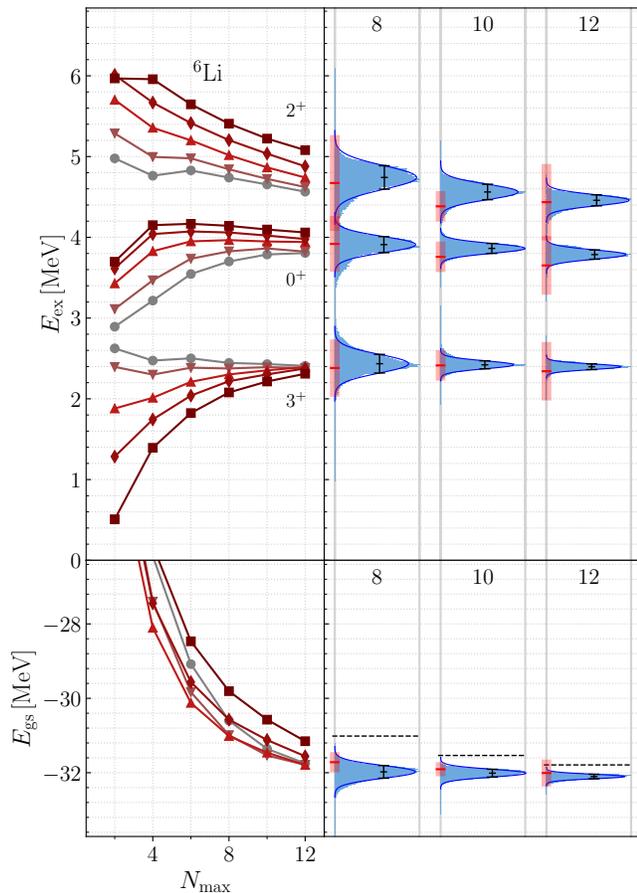


FIG. 6. Input data and network predictions for the ground state and excitation energy of ${}^6\text{Li}$. The lower left hand panels shows the NCSM input data for the ground-state prediction, while the upper panel illustrated the convergence behavior of the excited states, with data for $\hbar\Omega = 14, 16, 20, 24,$ and 28 MeV (gray to red). The right hand panels show the corresponding network predictions, together with the classical extrapolation results and the variational minimum for the ground-state energy (dashed line). Note that the ANN predictions for the excitation energies are based on the sample-wise difference of absolute energy predictions. The corresponding convergence patterns in the top left panel are only shown for comparison.

observables are depicted as a red bar, with the corresponding uncertainty as a red band around it.

One major observation we want to highlight is that the ANN predictions for the ground-state energies are remarkably stable with increasing model-space sizes and generally significantly outperform the classical extrapolations at $\mathcal{N}_{\text{max}} = 8$ across all considered nuclei when compared with the values obtained for large \mathcal{N}_{max} . The classical extrapolations are more susceptible to a systematic downwards trend with increasing model-space size, which has its roots in the fact that the convergence of ground-state energies is not quite exponential. The networks do not suffer from such a systematic, which facilitates a robust and accurate prediction of the converged energy in smaller model spaces.

Excitation energies are slightly more difficult, since the cor-

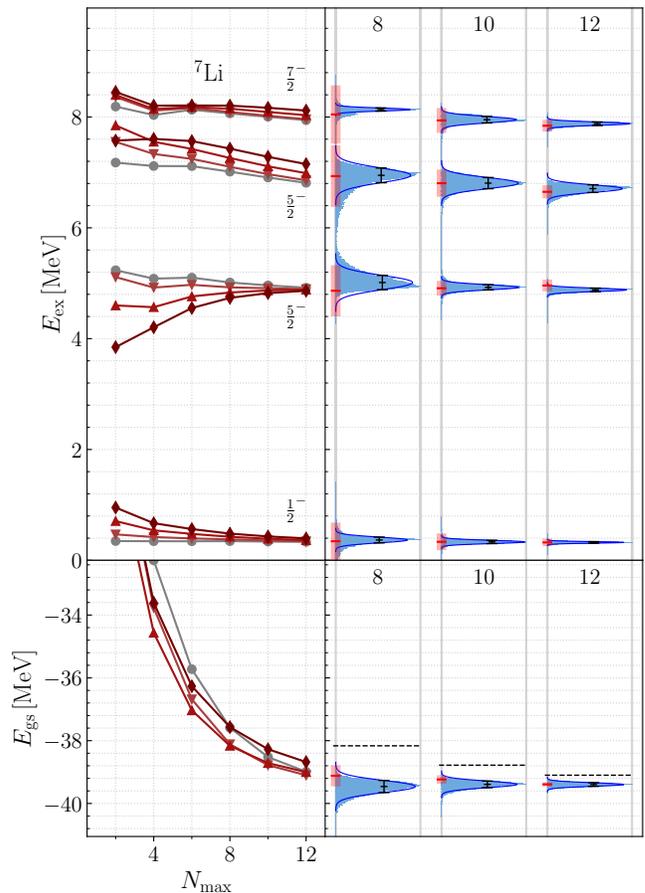


FIG. 7. Input data and network predictions for the ground state and excitation energy of ${}^7\text{Li}$ with $\hbar\Omega = 14, 16, 20,$ and 24 MeV (gray to red), analogous to Fig. 6

relation between the convergence of the ground and the excited state are relevant. Furthermore, excited-state energies, especially for more loosely bound states, converge slower than the ground-state energy. Together, this makes the prediction of the converged value much more challenging, especially at $\mathcal{N}_{\text{max}} = 8$. Nevertheless, the ANN results for excitation energies are generally very robust, in particular for the lowest excited states. Higher-lying, more loosely bound excited states like the 2^+ state in ${}^6\text{Li}$ or the $\frac{7}{2}^-$ state in ${}^7\text{Li}$ require $\mathcal{N}_{\text{max}} = 10$ to give accurate results. However, in these cases, the convergence behavior of the sequences changes significantly within the lower \mathcal{N}_{max} data points, so a slight influence on the ANN predictions is expected. For this reason, we recommend checking for the robustness of the predictions by comparing with predictions for other \mathcal{N}_{max} if available.

VIII. RESULTS FOR RADII

As already mentioned, the extrapolation of NCSM sequences for rms radii is very challenging, as their convergence behavior is far less constrained than for the ground-state en-

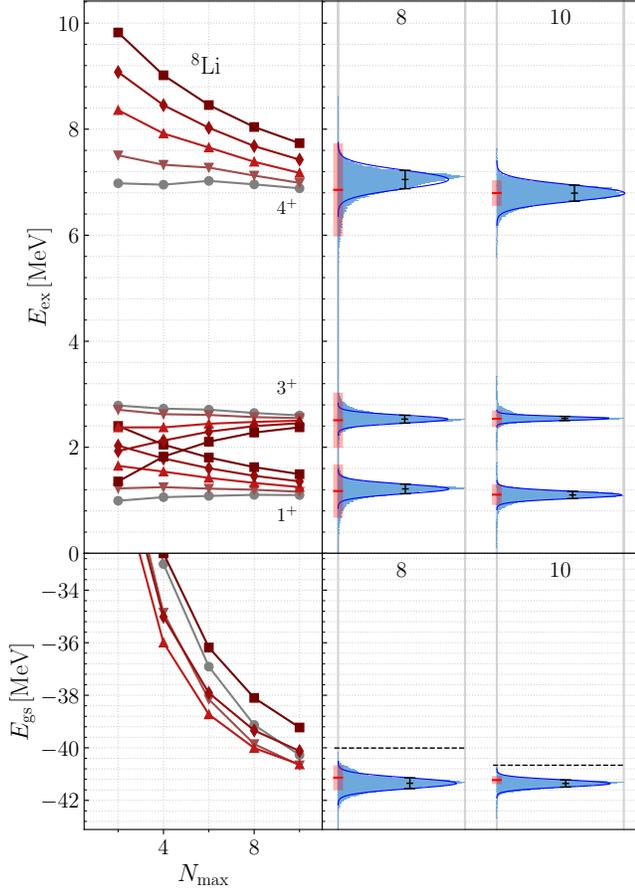


FIG. 8. Input data and network predictions for the ground state and excitation energy of ${}^8\text{Li}$ with $\hbar\Omega = 14, 16, 20, 24,$ and 28 MeV (gray to red), analogous to Fig. 6.

ergy. As a result a larger variety of convergence patterns can be observed and the networks have to learn to cope with the different patterns. Obviously, we have to train a new set of networks to pick up the convergence patterns inherent to the radius. We use the same network topology and input mode (MINMAX) as for the ground-state energy training, but train on data for mass rms radii instead. The evaluation step is performed analogous to the ground-state energies.

The application of the radius networks to ${}^2\text{H}$, ${}^3\text{H}$, and ${}^4\text{He}$ is illustrated in Fig. 10 and the final ANN predictions are summarized in Table I. The NCSM data for the mass rms radius, which is fed into the networks, is shown in the left-hand panels, while the ANN predictions are depicted on the right side, together with the fully converged values (green lines). As before, we use the semi-local NN+3N interaction at N^2LO with $\Lambda = 450$ MeV [25, 33] and $\alpha = 0.08$ fm 4 for this evaluation, which is not included in the training set.

We first observe that the convergence patterns for the radii are really very different from the ground-state energies. The radius is not constrained via the variational principle, which is evident from the fact that convergence is not monotonous. Depending on the selection of HO frequencies, whole sequences can converge from or from below. The convergence itself is

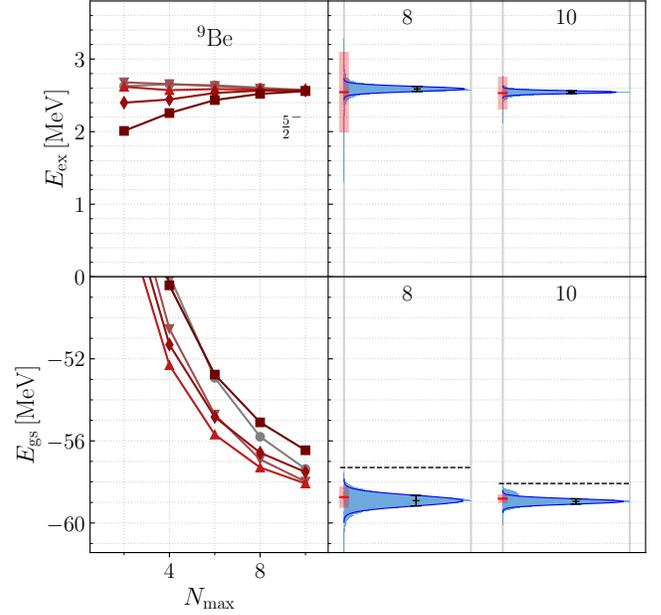


FIG. 9. Input data and network predictions for the ground state and excitation energy of ${}^9\text{Be}$ with $\hbar\Omega = 14, 16, 20, 24,$ and 28 MeV (gray to red), analogous to Fig. 6.

also significantly slower compared to energies, and some sequences show a jagged pattern, particularly at N_{max} values.

Despite of these characteristics, the ANN predictions for the ${}^2\text{H}$ and ${}^4\text{He}$ mass rms radius are in very good agreement with the fully converged value, even at $N_{\text{max}} = 8$. These two nuclei also demonstrate the two extremes of possible convergence patterns: With a selection of HO frequencies around the variational minimum for the ground-state energy, ${}^4\text{He}$ exhibits a very symmetric convergence, with some sequences converging from above and some from below. ${}^2\text{H}$ on the other hand exhibits convergence exclusively from below, which is the typical case for most nuclei we looked at. This reflects the distribution of target values for all training samples mentioned earlier, and depicted in Fig. 2. The predictions for ${}^3\text{H}$ show a light drift, particularly the $N_{\text{max}} = 8$ is slightly above the exact radius, which might be related to the strong zig-zag pattern for small N_{max} values in the input data.

To assess the performance of the networks beyond of the few-body systems used for training, we take a look at four p-shell nuclei. Figure 11 shows input data and ANN predictions for ${}^6\text{Li}$, ${}^7\text{Li}$, ${}^8\text{Li}$, and ${}^9\text{Be}$. The final predictions are summarized in Table I. This data again is extracted from the NCSM calculations performed for Ref. [34], which focused entirely on energies. Therefore, the selection of HO frequencies available is optimized for the description of energies and not for rms radii. As a result, the sequences predominantly converge from below and, particularly for ${}^6\text{Li}$ and ${}^7\text{Li}$, do not cover the range around the optimum frequency for the radius.

Despite these deficiencies in the evaluation data, the predictions are very stable across the board. For ${}^6\text{Li}$ and ${}^7\text{Li}$ a slight trend towards larger radii, which is compatible with the depicted uncertainties, can be observed with increasing N_{max} .

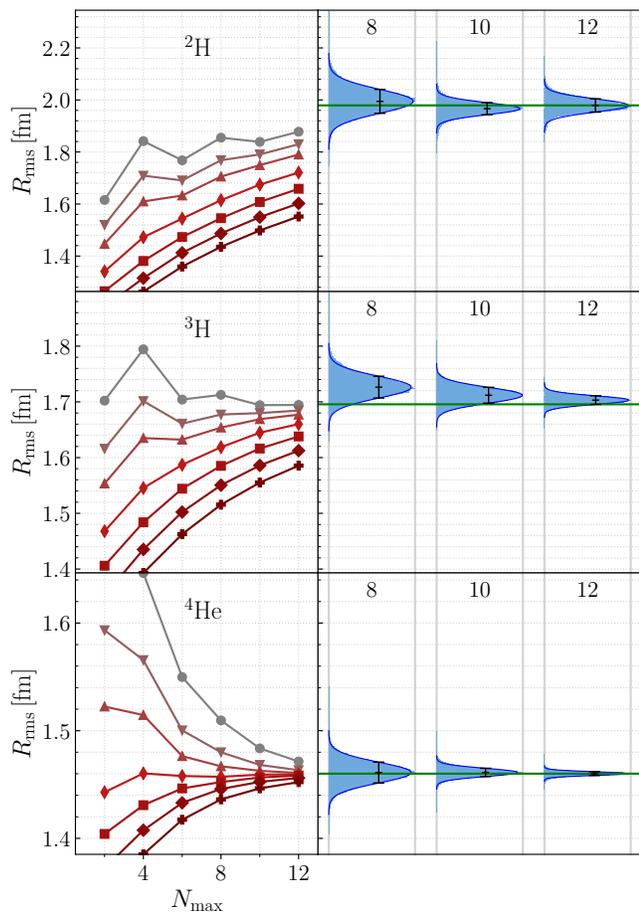


FIG. 10. Mass rms radii input data and resulting network predictions for few-body systems. The left hand panels show the NCSM data for $\hbar\Omega = 12, 14, 16, 20, 24, 28,$ and 32 MeV (gray to red), while the right hand panels show the ANN predictions for $N_{\max} = 8, 10,$ and 12 , with the green line indicating the fully converged value obtained at a large N_{\max} .

This shift gets smaller for larger model spaces, and is less pronounced for ${}^8\text{Li}$ and ${}^9\text{Be}$, where the chosen HO frequencies include more optimal values producing a flat convergence curve. Therefore, we expect that an optimized selection of HO frequencies in the NCSM calculations, leading to a more symmetric convergence pattern, will improve the predictions and alleviate the observed slight upwards trend. Another way to tackle such a systematic might be to restrict the training data set to sequences that converge exclusively from below, thereby, obtaining networks that are specialized for this situation. However, even with a non-optimized frequency selection and the general set of training data, the ANNs provide robust predictions for the converged rms radii with realistic data-driven uncertainties based on a rather limited set of evaluation data.

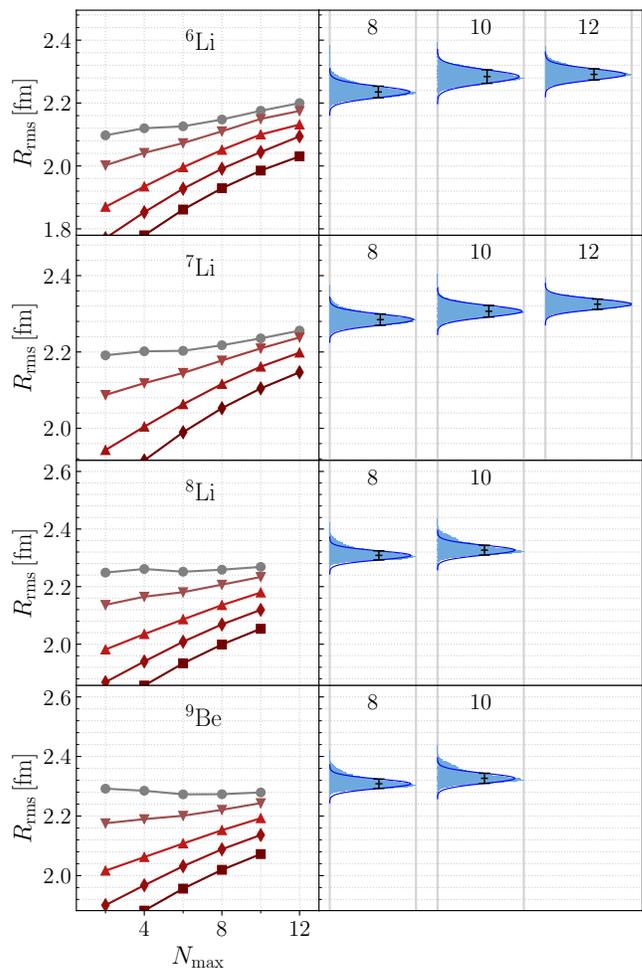


FIG. 11. Mass rms radius input data and network predictions for light p-shell nuclei with $\hbar\Omega = 14, 16, 20, 24,$ and 28 MeV (gray to red, ${}^7\text{Li}$ without $\hbar\Omega = 28$ MeV), analogous to Fig. 10.

IX. CONCLUSIONS

We have extended our ANN architecture to high-precision predictions of ground-state and excited-state energies and radii from non-converged NCSM calculations. The core idea of the framework is to learn convergence patterns from a large pool of training data obtained from few-body systems such as ${}^2\text{H}$, ${}^3\text{H}$, and ${}^4\text{He}$ with a range of different interactions, and apply the trained networks to NCSM data for a larger range of nuclei. This includes a statistical evaluation scheme, where many networks are trained and evaluated to obtain a distribution of predictions that allows for the extraction of statistically meaningful uncertainties. With the new MINMAX input mode we are able to extend this capability to radii as well, a much less constrained observable compared to energies. Additionally, we find improvements in accuracy coupled with more realistic uncertainties in the case of energies when compared with our previous generation of networks.

Based on these advances we are convinced that the ANN architecture is capable of handling other observable as well,

		$N_{\max} = 8$	10	12	conv.
^2H					
E_{gs}	ANN	-2.195(47)	-2.223(27)	-2.195(14)	-2.200
E_{gs}	extr.	-2.013(106)	-2.147(*)	-2.183(106)	-2.200
R_{rms}	ANN	1.995(46)	1.967(23)	1.979(26)	1.979
^3H					
E_{gs}	ANN	-8.482(65)	-8.509(33)	-8.487(12)	-8.481
E_{gs}	extr.	-8.434(110)	-8.480(109)	-8.473(24)	-8.481
R_{rms}	ANN	1.726(20)	1.712(14)	1.703(8)	1.696
^4He					
E_{gs}	ANN	-28.513(71)	-28.526(25)	-28.526(8)	-28.524
E_{gs}	extr.	-28.503(48)	-28.528(24)	-28.523(9)	-28.524
R_{rms}	ANN	1.461(10)	1.461(4)	1.460(2)	1.460
^6Li					
E_{gs}	ANN	-31.98(17)	-32.01(11)	-32.10(6)	-
E_{gs}	extr.	-31.72(27)	-31.91(19)	-32.01(36)	-
$E_{\text{ex}}(3^+)$	ANN	2.44(12)	2.42(5)	2.40(4)	-
$E_{\text{ex}}(3^+)$	extr.	2.38(35)	2.41(19)	2.34(36)	-
$E_{\text{ex}}(0^+)$	ANN	3.91(10)	3.86(6)	3.79(6)	-
$E_{\text{ex}}(0^+)$	extr.	3.92(35)	3.76(19)	3.65(36)	-
$E_{\text{ex}}(2^+)$	ANN	4.74(15)	4.56(10)	4.46(7)	-
$E_{\text{ex}}(2^+)$	extr.	4.67(59)	4.38(19)	4.44(47)	-
R_{rms}	ANN	2.235(18)	2.284(22)	2.291(18)	-
^7Li					
E_{gs}	ANN	-39.46(19)	-39.39(11)	-39.40(6)	-
E_{gs}	extr.	-39.12(34)	-39.24(13)	-39.39(7)	-
$E_{\text{ex}}(1/2^-)$	ANN	0.37(5)	0.33(3)	0.32(1)	-
$E_{\text{ex}}(1/2^-)$	extr.	0.34(34)	0.33(15)	0.32(7)	-
$E_{\text{ex}}(5/2^-)$	ANN	5.01(13)	4.93(5)	4.88(3)	-
$E_{\text{ex}}(5/2^-)$	extr.	4.86(46)	4.91(13)	4.96(11)	-
$E_{\text{ex}}(5/2^-)$	ANN	6.95(13)	6.81(10)	6.71(7)	-
$E_{\text{ex}}(5/2^-)$	extr.	6.93(55)	6.81(25)	6.65(12)	-
$E_{\text{ex}}(7/2^-)$	ANN	8.14(3)	7.95(6)	7.88(3)	-
$E_{\text{ex}}(7/2^-)$	extr.	8.05(53)	7.94(23)	7.84(11)	-
R_{rms}	ANN	2.285(15)	2.307(15)	2.325(14)	-
^8Li					
E_{gs}	ANN	-41.35(21)	-41.36(14)	-	-
E_{gs}	extr.	-41.14(47)	-41.23(16)	-	-
$E_{\text{ex}}(1^+)$	ANN	1.22(9)	1.10(7)	-	-
$E_{\text{ex}}(1^+)$	extr.	1.18(50)	1.11(19)	-	-
$E_{\text{ex}}(3^+)$	ANN	2.53(8)	2.54(4)	-	-
$E_{\text{ex}}(3^+)$	extr.	2.51(52)	2.54(16)	-	-
$E_{\text{ex}}(4^+)$	ANN	7.05(17)	6.79(15)	-	-
$E_{\text{ex}}(4^+)$	extr.	6.86(88)	6.80(24)	-	-
R_{rms}	ANN	2.308(16)	2.327(17)	-	-
^9Be					
E_{gs}	ANN	-58.91(26)	-58.95(13)	-	-
E_{gs}	extr.	-58.74(52)	-58.82(21)	-	-
$E_{\text{ex}}(5/2^-)$	ANN	2.59(4)	2.54(2)	-	-
$E_{\text{ex}}(5/2^-)$	extr.	2.54(55)	2.53(23)	-	-
R_{rms}	ANN	2.309(14)	2.324(13)	-	-

TABLE I. ANN predictions of ground state and excitation energies [MeV] and the mass rms radius [fm] for all discussed nuclei. Classical extrapolation results for energies are given as a comparison. Where available, the fully converged result is given as well. Uncertainties denoted as (*) are unreasonably large due to a breakdown of the classical extrapolation method.

especially electromagnetic moments and transition strengths. For electric quadrupole observables we expect a correlation to the rms radius, which can be exploited in the evaluation and construction of the networks. Of course, for the training of new ANN for such observables we need a sufficient amount of training data in very light nuclei. This can become an issue, since many electromagnetic observables, particularly transition strengths, do not naturally appear in bound few-nucleon system. To remedy this issue, we are currently looking into synthetic nuclei which emerge from modified Hamiltonians. We can, for example, enhance the interaction to create additional bound states in these small systems, making more complex observables accessible in few-body systems and to create a large library of training data.

ACKNOWLEDGMENTS

We thank Pieter Maris for providing us with the raw NCSM output for the calculations presented in Refs. [25, 34], which were performed at the Argonne Leadership Computing Facility (ALCF), a US Department of Energy Office of Science user facility, supported under Contract No. DE-AC02-06CH11357, using computing resources provided under the INCITE award ‘Nuclear Structure and Nuclear Reactions’ from the US Department of Energy, Office of Advanced Scientific Computing Research. All other numerical calculations have been performed on the LICHTENBERG II cluster at the computing center of the TU Darmstadt. This work is supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through the DFG Sonderforschungsbereich SFB 1245 (Project-ID 279384907) and the BMBF through Verbundprojekt 05P2021 (ErUM-FSP T07, Contract No. 05P21RDFNB).

- [1] B. R. Barrett, P. Navrátil, and J. P. Vary, Ab initio no core shell model, *Prog. Part. Nucl. Phys.* **69**, 131 (2013).
[2] P. Navrátil, S. Quaglioni, I. Stetcu, and B. R. Barrett, Recent developments in no-core shell-model calculations, *J. Phys. G:*

Nucl. Part. Phys. **36**, 083101 (2009).

- [3] R. Roth, Importance truncation for large-scale configuration interaction approaches, *Phys. Rev. C* **79**, 064324 (2009).
[4] D. C. Zheng, B. R. Barrett, L. Jaqua, J. P. Vary, and R. J. Mc-

- Carthy, Microscopic calculations of the spectra of light nuclei, *Phys. Rev. C* **48**, 1083 (1993).
- [5] K. Kowalski, D. J. Dean, M. Hjorth-Jensen, T. Papenbrock, and P. Piecuch, Coupled Cluster Calculations of Ground and Excited States of Nuclei, *Phys. Rev. Lett.* **92**, 132501 (2004).
- [6] W. Dickhoff and C. Barbieri, Self-consistent Green's function method for nuclei and nuclear matter, *Prog. Part. Nucl. Phys.* **52**, 377 (2004).
- [7] H. Hergert, S. Bogner, T. Morris, A. Schwenk, and K. Tsukiyama, The In-Medium Similarity Renormalization Group: A novel ab initio method for nuclei, *Phys. Rep.* **621**, 165 (2016).
- [8] J. Carlson, S. Gandolfi, F. Pederiva, S. C. Pieper, R. Schiavilla, K. E. Schmidt, and R. B. Wiringa, Quantum Monte Carlo methods for nuclear physics, *Rev. Mod. Phys.* **87**, 1067 (2015).
- [9] J. W. Clark, T. Lindenau, and M. L. Ristig, eds., *Scientific Applications of Neural Nets*, Lecture Notes in Physics, Vol. 522 (Springer Berlin Heidelberg, 1999).
- [10] A. Boehnlein *et al.*, *Colloquium* : Machine learning in nuclear physics, *Rev. Mod. Phys.* **94**, 031003 (2022).
- [11] J. Keeble and A. Rios, Machine learning the deuteron, *Phys. Lett. B* **809**, 135743 (2020).
- [12] C. Adams, G. Carleo, A. Lovato, and N. Rocco, Variational Monte Carlo Calculations of $A \leq 4$ Nuclei with an Artificial Neural-Network Correlator Ansatz, *Phys. Rev. Lett.* **127**, 022502 (2021).
- [13] A. Gnech, C. Adams, N. Brawand, G. Carleo, A. Lovato, and N. Rocco, Nuclei with Up to $A = 6$ Nucleons with Artificial Neural Network Wave Functions, *Few-Body Syst.* **63**, 7 (2022).
- [14] A. Lovato, C. Adams, G. Carleo, and N. Rocco, Hidden-nucleons neural-network quantum states for the nuclear many-body problem, *Phys. Rev. Research* **4**, 043178 (2022).
- [15] S. Athanassopoulos, E. Mavrommatis, K. Gernoth, and J. Clark, Nuclear mass systematics using neural networks, *Nucl. Phys. A* **743**, 222 (2004).
- [16] S. Akkoyun, T. Bayram, S. O. Kara, and A. Sinan, An artificial neural network application on nuclear charge radii, *J. Phys. G: Nucl. Part. Phys.* **40**, 055106 (2013).
- [17] G. A. Negroita *et al.*, Deep learning: Extrapolation tool for *ab initio* nuclear theory, *Phys. Rev. C* **99**, 054308 (2019).
- [18] W. G. Jiang, G. Hagen, and T. Papenbrock, Extrapolation of nuclear structure observables with artificial neural networks, *Phys. Rev. C* **100**, 054326 (2019).
- [19] M. Knöll, T. Wolfgruber, M. L. Agel, C. Wenz, and R. Roth, Machine learning for the prediction of converged energies from ab initio nuclear structure calculations, *Phys. Lett. B* **839**, 137781 (2023).
- [20] P. Maris, J. P. Vary, and A. M. Shirokov, *Ab Initio* no-core full configuration calculations of light nuclei, *Phys. Rev. C* **79**, 014308 (2009).
- [21] S. Bogner, R. Furnstahl, P. Maris, R. Perry, A. Schwenk, and J. Vary, Convergence in the no-core shell model with low-momentum two-nucleon interactions, *Nucl. Phys. A* **801**, 21 (2008).
- [22] R. Roth, J. Langhammer, A. Calci, S. Binder, and P. Navrátil, Similarity-Transformed Chiral $NN + 3N$ Interactions for the *Ab Initio* Description of C 12 and O 16, *Phys. Rev. Lett.* **107**, 072501 (2011).
- [23] R. Roth, A. Calci, J. Langhammer, and S. Binder, Evolved chiral $NN + 3N$ Hamiltonians for *ab initio* nuclear structure calculations, *Phys. Rev. C* **90**, 024325 (2014).
- [24] E. D. Jurgenson, P. Maris, R. J. Furnstahl, P. Navrátil, W. E. Ormand, and J. P. Vary, Structure of p-shell nuclei using three-nucleon interactions evolved with the similarity renormalization group, *Phys. Rev. C* **87**, 054312 (2013).
- [25] P. Maris *et al.* (LENPIC Collaboration), Light nuclei with semilocal momentum-space regularized chiral interactions up to third order, *Phys. Rev. C* **103**, 054001 (2021).
- [26] P. Navrátil, G. P. Kamuntavičius, and B. R. Barrett, Few-nucleon systems in a translationally invariant harmonic oscillator basis, *Phys. Rev. C* **61**, 044001 (2000).
- [27] V. Nair and G. E. Hinton, Rectified Linear Units Improve Restricted Boltzmann Machines, in *Proceedings of the 27th International Conference on Machine Learning* (Omnipress, 2010) pp. 807–814.
- [28] I. Loshchilov and F. Hutter, Decoupled weight decay regularization (2017), [arXiv:1711.05101](https://arxiv.org/abs/1711.05101).
- [29] K. A. Wendt, C. Forssén, T. Papenbrock, and D. Sääf, Infrared length scale and extrapolations for the no-core shell model, *Phys. Rev. C* **91**, 061301 (2015).
- [30] C. Forssén, B. D. Carlsson, H. T. Johansson, D. Sääf, A. Bansal, G. Hagen, and T. Papenbrock, Large-scale exact diagonalizations reveal low-momentum scales of nuclei, *Phys. Rev. C* **97**, 034328 (2018).
- [31] D. R. Entem, R. Machleidt, and Y. Nosyk, High-quality two-nucleon potentials up to fifth order of the chiral expansion, *Phys. Rev. C* **96**, 024004 (2017).
- [32] T. Hübner, K. Vobig, K. Hebeler, R. Machleidt, and R. Roth, Family of chiral two- plus three-nucleon interactions for accurate nuclear structure studies, *Phys. Lett. B* **808**, 135651 (2020).
- [33] P. Reinert, H. Krebs, and E. Epelbaum, Semilocal momentum-space regularized chiral two-nucleon potentials up to fifth order, *Eur. Phys. J. A* **54**, 86 (2018).
- [34] P. Maris, H. Le, A. Nogga, R. Roth, and J. P. Vary, Uncertainties in ab initio nuclear structure calculations with chiral interactions, *Front. Phys.* **11**, 1098262 (2023).
- [35] P. Maris, private communication.