

Machine-learning approach to finite-size effects in systems with strongly interacting fermionsNawar Ismail and Alexandros Gezerlis *Department of Physics, University of Guelph, Guelph, Ontario N1G 2W1, Canada*

(Received 7 July 2021; revised 8 September 2021; accepted 28 September 2021; published 5 November 2021)

We investigate the applicability of machine learning techniques in studying the finite-size effects associated with many-body physics. These techniques have an emerging presence in many-body theory as they have been used for interpolations, extrapolations, and in modeling wave functions. We will resolve several issues associated with machine learning and many-body calculations such as small datasets, outliers, and discontinuities, for the purpose of extrapolating finite calculations to macroscopic scales. We carry out a systematic investigation of two related systems by developing metrics that aim to avoid spurious effects and capture desired features. This work uses neural networks to extrapolate the unitary gas to the thermodynamic limit at zero range, which is otherwise difficult to reach. The effective mass of strongly interacting neutron matter is also studied and makes use of the noninteracting problem to resolve discontinuous predictions. For this investigation, we also carried out new auxiliary field diffusion Monte Carlo (AFDMC) calculations for a variety of densities and particle numbers. Ultimately, we demonstrate an effective utility for neural networks in this context.

DOI: [10.1103/PhysRevC.104.055802](https://doi.org/10.1103/PhysRevC.104.055802)**I. INTRODUCTION**

In recent years machine learning (ML) techniques have established themselves in quantum many-body theory as a set of essential and promising tools for solving a diverse range of novel and existing problems. In the study of nuclei, extrapolations from the no-core shell model along with coupled-cluster calculations have been carried out to determine nuclear masses and radii [1,2]. Machine learning has also been applied to improve computational efficiency [3–8] and to develop computationally feasible models, like those for wave functions [9–11]. In general, the utility provided by these machine learning models arises from their ability to capture underlying dependencies, allowing for data to be interpolated and extrapolated. Despite the ability of the networks to generalize, special attention must often be paid to the representation of the data through, for example, feature engineering, to avoid misleading extrapolations [12].

The *ab initio* study of strongly interacting systems is often limited by prohibitive computational costs, which constrain calculations to finite domains, typically in terms of particle number. In the study of the unitary gas (UG) and of neutron matter (NM), these limitations arise when solving the quantum many-body problem and prevent certain calculations from significantly exceeding ≈ 100 particles. A large focus is appropriately placed on investigating the difference between these finite calculations and the infinite system, which are known as the finite-size effects (FSE) [1,13–17]. For systems with a fixed density, the infinite system is referred to as the thermodynamic limit (TL) and best describes macroscopic matter. In general, these deviations tend to diminish as the number of particles rises, although the behavior is often neither smooth nor monotonic. Typically, careful analysis

is required to determine which finite particle number best matches the TL. The ability to interpolate these dependencies can provide great insight into the limiting behavior. Naturally, this type of problem falls in the domain of machine learning, and so we aim to apply those techniques to study these finite-size effects.

The neural network (NN) has become a ubiquitous machine learning model due to its effectiveness in a variety of problems. Such an approach takes in a set of inputs and, through some scheme, propagates it through the layers of the network until an output is achieved. The particular scheme used to derive an output depends on the organization and structure of the network, which is known as its architecture. The internal weights used in the model are typically organized into so-called hidden layers, which vary depending on the type of network. Feed-forward neural networks (FFNN) have been used extensively in nuclear physics [1,2,9,12,18]. Other effective neural networks are the Boltzmann machine (BM) [3,10,11,19] and recursive neural networks (RNN) [20]. Sometimes tailored architectures are required or are beneficial for solving particular problems [21].

As noted, an FFNN is organized into a set of layers, where the first layer receives input values that are propagated through the hidden layers until an output is produced [1,22]. Each layer is comprised of nodes that are connected to each node in the subsequent layer. These connections are viewed as weights since the value of a node, y , depends on the previous layer x through a weighted sum:

$$y = f(\mathbf{w} \cdot \mathbf{x} + b), \quad (1)$$

where \mathbf{w} and b are free parameters that can be tuned, and f is an activation function that can be used to introduce nonlinearities to the model. The process is iterated until the

final output(s) are generated. Multiple hidden layers provide a hierarchical structure to the network that can capture more complex features; however, training these networks can require more time and care to avoid problems such as gradient vanishing. However, one of the key properties of the neural network is the universal approximation theorem, which states that a single hidden layer is sufficient to approximate any continuous function, provided the layer has a large enough number of neurons [1,12]. Of course, it does not suggest the ideal number of nodes in a layer or what the values of the internal weights should be, and these must instead be found during an optimization procedure known as training. During this optimization procedure, the small datasets often encountered in nuclear physics can pose challenging issues.

To identify and resolve difficulties that may arise in this context, we will be applying machine learning techniques to two related systems. In the case of the unitary gas, we encounter outliers and determine their impact by studying how network predictions vary based on how much emphasis is placed on them during training. To mitigate the negative effects of a small dataset, the technique of data augmentation is used to increase the dataset size without significant additional cost [1]. Other issues arise during the study of the effective mass of strongly interacting neutron matter. We performed additional energy evaluations using the auxiliary field diffusion Monte Carlo (AFDMC) algorithm to generate a dataset, which is provided as Supplemental Material in Ref. [23]. Due to discontinuities in the dataset, the networks initially provided spurious predictions, which could be corrected by using the noninteracting problem. After systematically studying these effects, TL predictions were arrived at for both problems.

II. METHODS

A. Potentials

Neutron matter plays an important role in understanding neutron-rich nuclei and neutron stars (NS) [24,25] and it can be described by the following Hamiltonian:

$$\hat{H} = -\frac{\hbar^2}{2m} \sum_i \nabla_i^2 + \sum_{i<j} v_{ij} + \sum_{i<j<k} v_{ijk}, \quad (2)$$

which considers the kinetic energy, a two-body interaction, and a three-body interaction. Higher-order considerations, like four-body interactions, have an effect on the total energy of the many-body system that is an order of magnitude smaller than that arising from three-nucleon interactions [26]; thus, we may safely ignore them. Many forms of potential exist that capture varying aspects and levels of detail from the underlying nuclear interaction; typically either phenomenological [16,27–30] or effective-field theory potentials [17,17,31–40] are used.

In this work, high-precision calculations are carried out for the phenomenological two-body Argonne $v8'$ (AV8') and three-body Urbana IX (UIX) potentials. The Argonne potential comprises spin, tensor, spin-orbit, and isospin operators with radial dependencies that are tuned to a large body of neutron-proton scattering data resulting in high-quality fits

[16,24,30,41,42]. The Urbana potential is similarly fit to light nuclei and nuclear matter [16,43].

In the low-density regime of neutron star crusts, the potential can be effectively parametrized by the scattering length and effective range, which greatly simplifies many considerations [44,45]. Given the scale independence of the UG, the fine details of the interatomic potential have little impact on observables [46,47]. This universal behavior allows us to group energy calculations from multiple potentials. Here we employ the diffusion Monte Carlo (DMC) results reported on in Ref. [46] as the input; these include calculations from both the modified-Pöschl-Teller potential and the double Gaussian potential, which are given by

$$v_{PT}(r) = 4\mu_{PT}^2 \operatorname{sech}^2(\mu_{PT}r) \quad (3)$$

and

$$v_{2G}(r) = 3.144\mu_{2G}^2 (e^{-\mu_{2G}^2 r^2/4} - 4e^{-\mu_{2G}^2 r^2}), \quad (4)$$

respectively, where $\mu_{PT} = 2/r_e$, $\mu_{2G} = 3.952/r_e$, r_e is the effective range, and r is the distance between two particles. These potentials are essentially zero-range two-body s -wave interactions; there is no three-body interaction [46,48].

B. Quantum Monte Carlo

Energy calculations are carried out by using Eq. (2) for the ground-state energy through the use of quantum Monte Carlo (QMC) algorithms [17,29,30,36,49]. Auxiliary Diffusion Monte Carlo (AFDMC) is a specialization of QMC that allows for high-precision calculations for up to about 100 particles [50–52]. The computational complexity of this algorithm (with respect to particle number) is largely dominated by wave function evaluations and limits us to studying finite systems [25,41]. A periodic boundary condition is applied which acts to approximate the macroscopic scale. This boundary condition is applied on a cubic box of length L and results in the discretization of the available momentum states such that the allowed wave vectors are given by $\mathbf{k} = (2\pi/L)(n_x, n_y, n_z)$, and the n 's are restricted to integers. It is easy to see that there exist many combinations of n 's that lead to the same wave number/energy. At finite N , these calculations deviate from the (otherwise unknown) TL values. These deviations are referred to as the finite-size effects (FSE) and a systematic study is required to make claims about macroscopic neutron matter [13–16]. As in most QMC works, we here employ periodic boundary conditions; a more general scheme, twisted boundary conditions, leads to distinct behavior which would have to be separately modeled [53,54].

C. Machine learning

In a task known as *supervised learning*, machine learning algorithms learn to make predictions from datasets that contain labeled data [55]. This training process aims to minimize the deviations from the predictions of a model and the dataset. The goal here is that, in addition to reproducing the dataset, the machine learning model is also capable of extrapolations or interpolations. This property of generalization is essential to avoid *overfitting*, whereby the model fails to capture

underlying features and is only capable of reproducing the given examples [1,5,56].

For this work, the feed-forward neural network [1,2,55] will be used to perform various regressions. Incidentally, when we employ below the term “neural network” we typically refer to an *ensemble* of networks, generated by using random initial conditions. The internal weights of the network are tuned to minimize the distance between the predictions and the dataset. This tuning occurs during a training process that typically uses the back-propagation and gradient-descent algorithms (and its variants) [1,22]. Our machine-learning calculations will be carried out using the KERAS Python library with a TENSORFLOW back end and using densely connected layers.

A dataset consists of examples. Each example contains an input and an output, which may each consist of multiple values. Networks trained on a given dataset are susceptible to *overfitting*, which describes the tendency for networks to generate poor predictions for examples that are not included in the training dataset (despite having a low error on the training examples). To avoid this issue, datasets are typically divided into a *training* portion and *testing* portion. The training portion is used to optimize the internal weights, while the testing portion is used to measure the ability of the network to generalize since the testing dataset contains novel examples. A sophisticated version of this, known as *k*-fold cross-validation [57], can be used to improve data usage during training as is done in the UG investigation. This technique divides the dataset into *k* subsets; *k* - 1 subsets are used during training, while 1 subset is left out as the test set to evaluate model performance. Since this test set can be any of the *k* subsets, the results are averaged across all permutations. Meanwhile, the effective mass will require additional considerations.

The optimization procedure attempts to determine optimal values for the internal weights of a model that allows it to generalize well. In addition to the internal weights of the neural network, there are a set of *hyperparameters* that also impact the predictions of the trained network. These may include factors such as the size of the network, the number of training iterations (epochs), and the use of a regularizer, among others. The performance of networks with different hyperparameters can be measured according to a prescribed metric, like the cross-validation error. Networks that perform well here are typically selected as those that have captured the underlying features.

III. LEARNING THE UNITARY GAS

A. Dataset

The UG inputs we are faced with are parametrized by two parameters: the effective range r_e and a finite particle number N [46]. The effective range is typically expressed as a dimensionless quantity, $k_F r_e$, where k_F is the Fermi momentum. Similarly, the energy E is also represented as a dimensionless quantity $\xi = E/E_{FG}$, where E_{FG} is the energy of the free Fermi gas.

The dataset shown in Fig. 1 demonstrates that the effective-range dependence is strongly linear while the N dependence




FIG. 1. The unitary gas dataset plotted along with three types of linear fits given by Eq. (5), (6), and (7), which correspond to the “individual,” “all,” and “average” fits, respectively. The effective-range dependence is strongly linear while the particle-number dependence is more complex. At $N = 70$ and $N = 80$, the slopes of the different fits notably disagree almost exclusively at these points, which suggests that they may be spurious.

is more complex, as noted in Ref. [46]. This is useful since linear fits can be used to *augment* the dataset [1], by providing additional points of data without much cost. Although higher-order fits (e.g., quadratic) may have some benefit, many of the N only correspond to two data points, and so for consistency we will employ linear fits. The slopes and intercepts can be determined in multiple ways:

$$\xi_{\text{individual}}(k_F r_e; N) = S(N)k_F r_e + b(N), \quad (5)$$

$$\xi_{\text{overall}}(k_F r_e; N) = \langle S(N) \rangle k_F r_e + b(N), \quad (6)$$

$$\xi_{\text{average}}(k_F r_e; N) = \frac{S(N) + \langle S(N) \rangle}{2} k_F r_e + b(N), \quad (7)$$

where $S(N)$ and $b(N)$ refer to the slope and intercept of the line that best fits through the points at a given N . The angle bracket notation denotes the average slope across all N . Studying these different fits will provide insight into the dataset. The individual fit captures only local effects, the overall captures only global effects, and the average attempts to capture both. These linear fits are shown in Fig. 1. Upon comparing these, outliers at $N = 70$ and $N = 80$ are identified. Since we are aiming to capture large N , removing these outliers prematurely may be counterproductive, so determining their impact is important.

B. Outliers

These outliers occur particularly at large N . To understand their effect, we train networks on datasets that emphasize different N through the means of upsampling. Let us denote the number of data points at a given particle number in the original dataset as C_N^{original} . We can tune the balance by linearly interpolating according to

$$C_N^{\text{balanced}}(t) = t C_N^{\text{original}} + (1 - t) \max(\{C_i^{\text{original}}\}). \quad (8)$$




FIG. 2. The dependence of the UG predictions on the dataset weightings is shown for three different datasets. These datasets consider all points included (left panel), an arbitrary pair removed (central panel), and the outliers removed (right panel). The TL energy estimates at 0 range are shown as a function of the dataset weighting t . In general, the estimates agree within a smaller margin when t corresponds to the original dataset (emphasis on low N), but diverge when t corresponds to the reflected dataset (emphasis on larger N and consequently the skewed slopes associated with the outliers).

where the second term picks the maximum number of data points across particle numbers. There are a few key values: $t = 1$ describes the original dataset, $t = 0$ describes a uniform distribution, and $t = -1$ describes a *reflected* distribution. The total counts here depend on t , which is undesired, and so the counts are normalized according to

$$C_N^{\text{fixed}}(t) = \left\lceil \frac{F}{\sum_N C_N^{\text{balanced}}(t)} \right\rceil C_N^{\text{balanced}}(t), \quad (9)$$

where the ceiling function is used to provide an integer count at least as large as F ; this F is the hyperparameter controlling the number of data points.

To isolate the effect of the outliers, we compare three datasets that consider all N , removing the outliers at $N = 70$, and $N = 80$, and removing an arbitrary pair (as a control). When the distribution favors low N (near $t = 1$), Fig. 2 shows that the predictions agree regardless of the fit used. However, as the distribution shifts more emphasis towards the outliers (near $t = -1$), there is a significant deviation between different fits. This occurs very strongly in the all-in and arbitrarily removed datasets, but not when the outliers are removed. The rightmost panel removes the suspected outliers resulting in a significantly reduced variance in the estimates. By contrast, removing a similar but arbitrary pair (50, 90) does not have this effect. This suggests that this reduced variance is due to removing the outliers, and not just removing arbitrary points. This suggests two things: the inclusion of the outliers *does* skew the predictions when emphasis is placed on them, and the original distribution ($t = 1$) does not provide a sufficient emphasis on the high N to capture the associated effects. It is therefore appropriate to remove the outliers and use a uniform distribution for training.

C. Removing pairs

So far, we have only validated this effect against a single arbitrary pair being removed. To solidify this, we will generalize this procedure to multiple other pairs. For a fixed F , the datasets have three parameters: the linear fit form, the removed pair, and the balance parameter t . To measure the variation in predictions due to the different fits s , we take the difference between the maximum and minimum predictions for a network trained with a given pair \mathcal{N} removed,

$$s(\mathcal{N}) = \max_{t,f} [\xi_f(\mathcal{N}, t)] - \min_{t,f} [\xi_f(\mathcal{N}, t)], \quad (10)$$

where $\xi_f(\mathcal{N}, t)$ is an energy prediction, $t \in [-1, 1]$, and $f \in \{\text{individual, overall, average}\}$ is the fit form used.

Unfortunately, this would require a large number of networks to be trained. To simplify this, we turn to Fig. 2 to construct a surrogate function which is easier to evaluate, but still captures the spread. Since the energy predictions tend to increase steadily from $t = 1$ to $t = -1$ we evaluate only at these anchor points, which loosely captures half of the total spread, now defined as

$$s(\mathcal{N}) \equiv \max_f [\xi_f(\mathcal{N}, -1)] - \min_f [\xi_f(\mathcal{N}, +1)]. \quad (11)$$

This only requires six networks to be trained per removed pair and is shown in Fig. 2.

To estimate the spread associated with a given particle number (instead of a pair), one of the particle numbers in the pair is fixed to N . Then, the spread is averaged over a set of values for the other particle numbers N^* ,

$$Z(N) \propto \sum_{N^*} s(\{N, N^*\}). \quad (12)$$

The influence on the spread for each particle number is shown in Fig. 3. At $N = 70$ and $N = 80$ the spread is significantly reduced. Thus we can claim that specifically removing the outliers has a unique effect and so they are removed.

D. Hyperparameter optimization and predictions

Having isolated and controlled small sample effects, we are now capable of carrying out a hyperparameter optimization. The individual fit will be used to augment the dataset since the other linear fits have served their purpose of identifying the outliers. The hyperparameters under consideration are training dataset size, number of epochs, and hidden layer size. (A single “epoch” is when the whole dataset has been used once; it is standard to randomly select individual samples from the entire dataset many times over, i.e., for multiple epochs.) The training dataset size corresponds to F , noting that the actual count $C_N^F(t)$ gets rounded up according to the ceiling in Eq. (9), and is then further split into a training and testing dataset.

Performing the grid search depicted in Fig. 4, we find that epochs between 5000 and 10 000, and hidden units around 50 have saturated performance.

The training size requires further investigation since it has confounding effects. Ultimately, as the number of interpolated points increases arbitrarily, the training and testing set end up being strongly correlated. This means that increasing the value




FIG. 3. To verify that the pair of outliers identified in Fig. 2 has a unique effect we study the removal of different pairs from the dataset. Networks are trained on different datasets, each with an arbitrary pair removed. The spread is averaged over predictions for pairs containing N , which produces the values shown in the figure, according to Eq. (12). There is a significant reduction at both $N = 70$ and $N = 80$, which suggests that the outliers do uniquely skew the predictions.

of F beyond a certain point no longer provides an independent test. With the other hyperparameters fixed, we can perform TL, 0-range predictions while varying the dataset size to identify signs of overfitting. This is done in Fig. 5 which identifies overfitting for $F \geq 5000$ and underfitting for $F \leq 1000$. This leaves the optimal network to have $F \approx 2000$.

The final step is to extract the TL, 0 range prediction from our optimal networks, which works out to be $\xi = 0.389(1)$. There have been many calculations for the UG due to its great importance to many-body physics. Calculations are generally carried out as upper-bound estimates which have steadily decreased over time [58–60]. Reference [46], whose raw data we employed as input for the present study, used simpler extrapolation techniques to arrive at the value 0.3897(4). The current leading theoretical value made use of a very large lattice calculation to arrive at 0.372(5) [15,58]. Earlier experiments




FIG. 4. The relative fivefold cross-validation error for various hyperparameters is shown, where darker color indicates poorer performance. There are three parameters under consideration: each subplot contains a different dataset size F , with the number of training epochs along the bottom, and the number of hidden units along the vertical. In all cases, the performance stops improving after about 50 hidden units. For dataset sizes that are sufficiently large, the performance plateaus after about 5000 epochs. The consistent improvements resulting from increasing the dataset size has a confounding effect, which is discussed in the main text.




FIG. 5. A set of predictions from optimized neural networks (solid bands) trained on different dataset sizes, F , ordered in the caption the same way as in the plot itself at $N = 100$. These are plotted alongside the original dataset (dotted) which is linearly extrapolated to 0 range. The outliers at $N = 70$ and $N = 80$ are hollow to denote that they were removed, as discussed in Sec. III D. The kink at $N \approx 140$ in the black curve (solid) is evidence of overfitting; meanwhile the inability of the $F = 1000$ curve (forward slash) to capture the peak at low N is evidence of underfitting. The prediction from $F = 2000$ (backward slash) is thus selected as the optimal network.

predicted values above 0.39 [61,62], but have recently been improved to 0.376(4) [63]. Through additional considerations for the uncertainty involved, this estimate was lowered even further to 0.370(5)(8) [64], in close agreement with the theoretical value. These calculations tend to predict lower energies than these network predictions, but this is not unexpected, given that DMC obeys a variational (i.e., upper bound) property. The networks do well to interpolate the provided dataset, which provides upper bound estimates on the true energy.

IV. LEARNING THE EFFECTIVE MASS

A. Dataset

The effective mass arises in Landau Fermi liquid theory (LFLT) [65] and can be directly related to observables like the specific heat of neutron matter [66]. It can be extracted by studying the dispersion relation for an excited particle in NM, as done in Ref. [67] and Ref. [25]. The quasiparticle energy $\Delta E_{TL}^{(k_{TL})}$ is related to the momentum k_{TL}^2 through the dispersion relation

$$\Delta E_{TL}^{(k_{TL})} = \frac{\hbar^2 k_{TL}^2}{2m^*}, \quad (13)$$

where m^* is the effective mass. The subscripts “TL” refer to the notion that these quantities make use of an extrapolation prescription derived in Ref. [67], which aims to reduce FSE. Despite this reduction, Eq. (13) still depends on the number of particles in our simulation, N , and so the FSE must still be studied.

At a given N , we consider a system in its ground state and the associated energy, E_N . To probe the dispersion relation, we also consider adding a particle to this system in an excited




FIG. 6. Interpolations predicted by the networks for the extrapolated quasiparticle energy $\Delta E_{TL}^{(k_{TL})}$ (top panel) and extrapolated momentum k_{TL}^2 (bottom panel) expressed in MeV, at a density of 0.05 fm^{-3} . Both networks are trained on the same domain for consistency, although we show more points for the momentum since it is analytic. In both panels, the predictions (dashed) and calculated values (dots, solid) for the first four excited states are shown, increasing from bottom to top. Both extrapolated quantities have a similar dependence, which is generally captured by the networks. As expected, the discontinuities are not captured. However, since both models experience this same bias, the discontinuous predictions are resolved in the fit calculation for the effective mass.

state of momentum k and the associated energy of the system, $E_{N+1}^{(k)}$. To access the TL behavior of these systems, the extrapolation prescription can be applied to both the energy and the momentum. The extrapolated momentum can be expressed in terms of single-particle states:

$$\begin{aligned} k_{TL}^2 &= k^2 - k_{F,N}^2 + k_F^2 \\ &= \left(2\pi\sqrt{\frac{n}{N}}\right)^2 \bar{n}^2 - \left(2\pi\sqrt{\frac{n}{N}}\right)^2 \bar{n}_o^2 + (3\pi^2 n)^{2/3}, \end{aligned} \quad (14)$$

where \bar{n}_o and \bar{n} are the integer momenta for the ground state and excited state, respectively, not to be confused with the particle number density, n . Certain momentum values are inaccessible, due to the periodic boundary condition imposed in QMC, resulting in the discontinuities found in Fig. 6. The quasiparticle energy is given in Ref. [25] as

$$\Delta E_N^{(k)} = E_{N+1}^{(k)} - E_N + \frac{2}{5}\xi E_F, \quad (15)$$

where ξ is the Bertsch parameter, and E_F is the Fermi energy. The extrapolated version of this, $\Delta E_{TL}^{(k_{TL})}$ can be expressed in terms of the extrapolated momentum, a constant offset, and a potential energy term:

$$\Delta E_{TL}^{(k_{TL})} = \Delta U_N^{(k)} + \frac{2}{5}(\xi - 1)E_F + \frac{\hbar^2 k_{TL}^2}{2m}, \quad (16)$$

where $\Delta U_N^{(k)}$ is the difference in potential energy between both systems. The FSE introduced by this energy term are not large since, in neutron matter, the kinetic energy tends to be the dominant FSE contribution due to the small effective range of the interaction. In Fig. 6, this is evident since the

difference from the momentum is small. The discontinuities common to both the energy and momentum are due to the discretized nature of the available momentum states.

The effective mass is extracted from a linear fit between these quantities, according to Eq. (13). At a given density and particle number, the dispersion is studied by considering different excited states. To probe near the Fermi surface, as done in Ref. [25], four excited states are used in the fit. The quasiparticle energy is an interacting quantity and is therefore limited by QMC to finite N . By contrast, the momentum is a noninteracting quantity and can be computed for arbitrarily large systems.

In the ML task, the mapping to learn for the energy is

$$(N, n) \xrightarrow{\text{energy networks}} \{\Delta E_{TL}^e\}_{e=1}^4, \quad (17)$$

where the superscript e labels the ordering of excited states, and n again corresponds to the density.

The machine learning task is to interpolate these calculations to extract the underlying particle-number dependency. This mapping naturally suggests a FFNN with two inputs and four outputs. Alternate schemes may consider breaking the mapping into individual excitations; however, this would lose shared information between the different excited states. This notion of *feature engineering*, whereby the structure of the inputs/outputs is selected with the aim of simplifying the model's learning task, is important to consider when performing extrapolations [12].

Since the momentum is a noninteracting quantity, it can be calculated to very large N . However, spurious effects are introduced in the fitting procedure if these are used. Networks tend to smooth the predictions and struggle with the discontinuities shown in Fig. 6. The training dataset for the quasiparticle energy contained evaluations for $N \in [1, 2, 3, 4, 5, 6, 7, 19, 27, 33, 45, 57, 70]$ and $n \in [0.02, 0.04, 0.06, 0.08, 0.1, 0.12, 0.14, 0.16, 0.18, 0.2]$. During the fitting procedure, the mismatch between the energy and momentum continuities results in spurious predictions, as shown in Fig. 7.

The solution here is to separately train networks to learn the momentum and join those results with the energy during the fit, as shown in Fig. 7. For the momentum, the mapping to learn is

$$(N, n) \xrightarrow{\text{momentum networks}} \{k_{TL}^e\}_{e=1}^4. \quad (18)$$

Fitting Eqs. (17) and (18) as per Eq. (13) provides an effective-mass prediction. Alternative formulations that attempted to resolve these discontinuities included the use of custom loss functions, and the use of transfer learning to initialize the networks with noninteracting trends. Ultimately, the momentum networks were required and sufficient. The hyperparameter optimization here is more difficult than in the UG case, since there are additional considerations for the two networks and the different features they may capture. For this, we will seek additional metrics.

B. Selecting similar networks

The discontinuous effective-mass predictions can be resolved by training networks to learn both the energy and




FIG. 7. The effective masses predicted by neural networks at a density of 0.05 fm^{-3} are contrasted for two fit methods. These predictions constitute an interpolation, since this density was not included during training. In the top panel, the energy predictions are fit to the analytic momentum, which results in discontinuities. In the bottom panel, the energy predictions are fit with predictions from a separate network that was trained on the momentum, which resolves the discontinuities. The smoothing effects shared by both network predictions eliminate the discontinuities in the fits. At low particle numbers, the effective-mass predictions are more sensitive to the large energy/momentum values (shown in Fig. 6) which results in more variability. At high N , the predictions begin struggling to extrapolate and must contend with the larger QMC error.

momentum N dependence, and combining the results in a fit. Since the energy and momentum networks are independent, different details may be captured between the two. While we expect the energy and momentum trends to differ according to Eq. (16), the ability of the networks to capture these details may vary depending on the hyperparameters associated with each network. So, both network types are trained under a variety of hyperparameters, but in contrast to the UG, the testing error is not a sufficient metric. Ultimately, the prediction quality depends on the combination of networks, whereas the testing error provides an independent metric. This can cause issues since two networks with similar testing errors may perform better/worse at different regions, resulting in a fit that could potentially be poor along the whole domain. To combat this, we aim to construct a selection criterion that identifies networks that capture similar features.

The energy and momentum dependencies experience discontinuities due to the discrete momentum states; however, the effective mass does not. In fact, the effective mass is relatively constant as shown in Fig. 7. The claim that the effective mass *is* constant would be too strong, however the effective-mass predictions should have a low variance across N . At a given density, we can evaluate the variance to define the *disagreement factor* D as a metric for the quality of a prediction:

$$D \equiv \text{Var} \left(\left\{ \frac{m^*}{m}(n, N) \right\}_{N=N_{\text{low}}}^{N_{\text{high}}} \right), \quad (19)$$




FIG. 8. The effective-mass predictions m^*/m as a function of the disagreement factor D for a density of 0.04 fm^{-3} (left) and 0.18 fm^{-3} (right) for networks with various hyperparameter settings. This contrasts with the effect of increasing density; the networks at a lower density are more variable than those at higher densities. This is likely due to the reduced energy scale, which requires increased precision. Regardless of the density, there is a common trend that the network predictions limit towards some value as the disagreement factor decreases. At each density, the bottom 10% (the rightmost portion) are selected to calculate an effective-mass statistic, as done for Fig. 9. These networks averaged their predictions over the domain $N \in [0, 130]$.

where n is the number density, and the variance is evaluated over particle numbers $N \in [N_{\text{low}}, N_{\text{high}}]$. This interval is used to exclude particle numbers where QMC and/or the networks may perform poorly. The predictions in the intermediate regions are generally stable, but may vary near the extremes. To identify possible boundary effects, we consider two domains: $N \in [0, 130]$, and $N \in [40, 120]$. The effective-mass predictions averaged over a given interval are used as a statistic. The effective-mass predictions from networks with lower disagreement factors have best captured the overall dependence and are therefore used to generate an estimate. In general, the effective-mass predictions tend to converge to some value as the disagreement factor increases as shown in Fig. 8. Each point on this plot corresponds to a pair of neural network ensembles and their associated set of hyperparameter values. The average prediction of networks with disagreement factors in the lowest 10% is chosen as a statistic. At lower densities, the network predictions are less consistent. The networks may be struggling due to the difference in scale between higher and lower densities.

C. Predictions for density dependence

We now have a systematic procedure for training neural networks to predict effective masses that are continuous. The use of both energy and momentum networks resolved the discontinuous predictions, and the disagreement factor was used to assess the effective-mass prediction quality. The

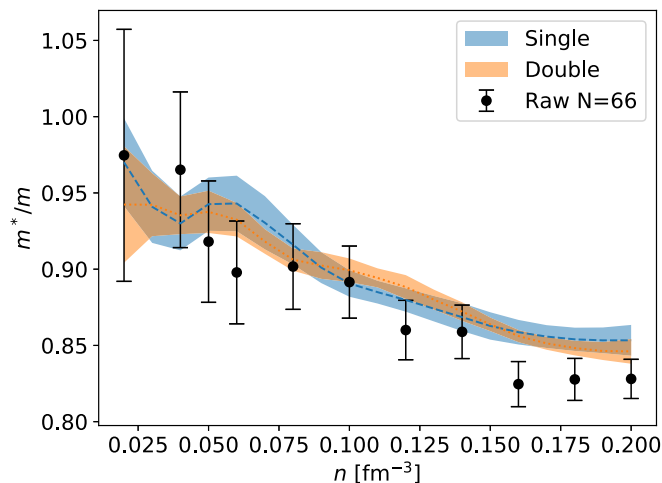


FIG. 9. Network predictions for the effective mass m^*/m as a function of number density n along with raw AFDMC calculations for $N = 66$. The blue dashed trend (which matches the raw calculation at the lowest density) shows predictions from networks with a single hidden layer, while the orange dotted trend uses networks with two equally sized hidden layers. These predictions were averaged over the domain $N \in [0, 130]$. Overall, the agreement between the networks is quite good.

predictions are averaged over a domain which is then used as an effective-mass statistic at a given density. This procedure can be repeated for a variety of densities as done in Fig. 9. We consider comparisons between different averaging domains and the effect of multiple hidden layers. In general, the agreement between all the predictions is quite good and they tend to agree within error.

Predictions at the lowest densities experience the highest variance. Networks trained on the larger domain ($N \in [0, 130]$) tend to capture the overall effects with less uncertainty than the smaller domain. Although a network with a single hidden layer can in theory approximate a function as well as a multilayered network, we would like to determine if multiple hidden layers may capture particular effects. To investigate this, we consider networks with two hidden layers of the same size. The predictions are similar, but tend to differ in subtle effects.

As a comparison to previous work, the effective-mass predictions of Ref. [25] take the AFDMC calculation at $N = 66$ as the estimate for each density. There is a good amount of overlap between the predictions and these raw estimates. In addition to the use of neural networks to carry out this regression analysis, another significant contrast is that this work makes use of many particle numbers to provide a density estimate. For these reasons, the networks are expected to have captured additional details that would otherwise be missed. The networks are able to interpolate between densities and particle numbers and so effects like the slightly higher effective masses predicted by the networks at higher

densities are expected to be better descriptions of the overall behavior.

V. SUMMARY & CONCLUSION

We trained feed-forward neural networks on two related nuclear systems to study finite-size effects through machine learning techniques to understand potential issues, demonstrate techniques to resolve them, and improve predictions. The difficulty of carrying out high-precision calculations through QMC restricts our datasets to small sizes. We found the technique of data augmentation effective to mitigate these issues, although some care was required to account for introduced correlations. Despite increasing the number of data points available, the accuracy of the model still depends on the quality of the original dataset. Outliers can skew the data, but due to the scarcity of data, we take care to demonstrate their negative impact before justifying their removal. After validating the dataset, a hyperparameter optimization was carried out by using k -fold cross validation. The optimized networks were able to produce TL estimates for the zero-range limit of the UG, that were not trivial to access directly in a QMC approach.

Networks can also struggle with other dataset effects, like discontinuities. They tend to predict smooth trends and struggle to capture sudden jumps. This smoothing bias introduces an inconsistency between the energy and momentum that was remedied by independently training networks to predict the associated momentum; as long as both quantities experienced similar behavior, the fitting process resulted in a continuous effective-mass prediction. The notion of *network disagreement* was introduced to identify networks that satisfied this property. Following this, the networks were trained on data resulting from raw AFDMC energy calculations, which are given as Supplemental Material in Ref. [23]. The TL density dependence of the effective mass for networks with single and double hidden layers were compared to raw calculations stemming from previous work.

Overall, we have found that neural networks provide a versatile tool that allows one to capture finite-size effects for both ground-state and excited-state properties. This is especially important when there is no *a priori* analytic expectation of what the N -dependence should be, as is often the case for strongly interacting systems. As seen in the results reported on in this work, it is possible to use machine learning to fold in the entire N dependence, without having to invoke simplifying approximations.

ACKNOWLEDGMENTS

This work was supported by the Natural Sciences and Engineering Research Council (NSERC) of Canada, the Canada Foundation for Innovation (CFI), and the Early Researcher Award (ERA) program of the Ontario Ministry of Research, Innovation and Science. Computational resources were provided by SHARCNET and NERSC.

- [1] W. G. Jiang, G. Hagen, and T. Papenbrock, *Phys. Rev. C* **100**, 054326 (2019).
- [2] G. A. Negroita, J. P. Vary, G. R. Luecke, P. Maris, A. M. Shirokov, I. J. Shin, Y. Kim, E. G. Ng, C. Yang, M. Lockner *et al.*, *Phys. Rev. C* **99**, 054308 (2019).
- [3] L. Huang and L. Wang, *Phys. Rev. B* **95**, 035105 (2017).
- [4] J. Liu, Y. Qi, Z. Y. Meng, and L. Fu, *Phys. Rev. B* **95**, 041101(R) (2017).
- [5] H. Shen, J. Liu, and L. Fu, *Phys. Rev. B* **97**, 205140 (2018).
- [6] J. Nelson, R. Tiwari, and S. Sanvito, *Phys. Rev. B* **99**, 075132 (2019).
- [7] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, *Nat. Phys.* **14**, 447 (2018).
- [8] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, *Nat. Mach. Intell.* **1**, 155 (2019).
- [9] R.-D. Lasserri, D. Regnier, J.-P. Ebran, and A. Penon, *Phys. Rev. Lett.* **124**, 162502 (2020).
- [10] G. Carleo and M. Troyer, *Science* **355**, 602 (2017).
- [11] X. Gao and L.-M. Duan, *Nat. Commun.* **8**, 662 (2017).
- [12] A. Pastore and M. Carmini, *J. Phys. G: Nucl. Part. Phys.* **48**, 084001 (2021).
- [13] H. Kwee, S. Zhang, and H. Krakauer, *Phys. Rev. Lett.* **100**, 126404 (2008).
- [14] S. Giorgini, L. P. Pitaevskii, and S. Stringari, *Rev. Mod. Phys.* **80**, 1215 (2008).
- [15] J. Carlson, S. Gandolfi, and A. Gezerlis, *Prog. Theor. Exp. Phys.* **2012**, 01A209 (2012).
- [16] S. Gandolfi, A. Y. Illarionov, K. E. Schmidt, F. Pederiva, and S. Fantoni, *Phys. Rev. C* **79**, 054005 (2009).
- [17] A. Gezerlis, I. Tews, E. Epelbaum, S. Gandolfi, K. Hebeler, A. Nogga, and A. Schwenk, *Phys. Rev. Lett.* **111**, 032501 (2013).
- [18] C. Adams, G. Carleo, A. Lovato, and N. Rocco, *Phys. Rev. Lett.* **127**, 022502 (2021).
- [19] E. Rrapaj and A. Roggero, *Phys. Rev. E* **103**, 013302 (2021).
- [20] D. Guest, K. Cranmer, and D. Whiteson, *Annu. Rev. Nucl. Part. Sci.* **68**, 161 (2018).
- [21] K. Raghavan, P. Balaprakash, A. Lovato, N. Rocco, and S. M. Wild, *Phys. Rev. C* **103**, 035502 (2021).
- [22] J. Keeble and A. Rios, *Phys. Lett. B* **809**, 135743 (2020).
- [23] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevC.104.055802> for details about the QMC calculations.
- [24] S. Gandolfi, A. Gezerlis, and J. Carlson, *Annu. Rev. Nucl. Part. Sci.* **65**, 303 (2015).
- [25] M. Buraczynski, N. Ismail, and A. Gezerlis, *Eur. Phys. J. A* **56**, 112 (2020).
- [26] I. Tews, T. Krüger, K. Hebeler, and A. Schwenk, *Phys. Rev. Lett.* **110**, 032504 (2013).
- [27] J. Carlson, J. Morales, V. R. Pandharipande, and D. G. Ravenhall, *Phys. Rev. C* **68**, 025802 (2003).
- [28] A. Gezerlis and J. Carlson, *Phys. Rev. C* **81**, 025803 (2010).
- [29] S. Gandolfi, J. Carlson, and S. Reddy, *Phys. Rev. C* **85**, 032801 (2012).
- [30] M. Baldo, A. Polls, A. Rios, H.-J. Schulze, and I. Vidaña, *Phys. Rev. C* **86**, 064001 (2012).
- [31] K. Hebeler and A. Schwenk, *Phys. Rev. C* **82**, 014314 (2010).
- [32] R. Machleidt and D. Entem, *Phys. Rep.* **503**, 1 (2011).
- [33] G. Hagen, T. Papenbrock, A. Ekström, K. A. Wendt, G. Baardsen, S. Gandolfi, M. Hjorth-Jensen, and C. J. Horowitz, *Phys. Rev. C* **89**, 014319 (2014).
- [34] A. Carbone, A. Rios, and A. Polls, *Phys. Rev. C* **90**, 054322 (2014).
- [35] A. Roggero, A. Mukherjee, and F. Pederiva, *Phys. Rev. Lett.* **112**, 221103 (2014).
- [36] G. Wlazłowski, J. W. Holt, S. Moroz, A. Bulgac, and K. J. Roche, *Phys. Rev. Lett.* **113**, 182503 (2014).
- [37] V. Somà, A. Cipollone, C. Barbieri, P. Navrátil, and T. Duguet, *Phys. Rev. C* **89**, 061301 (2014).
- [38] I. Tews, S. Gandolfi, A. Gezerlis, and A. Schwenk, *Phys. Rev. C* **93**, 024305 (2016).
- [39] M. Piarulli, A. Baroni, L. Girlanda, A. Kievsky, A. Lovato, E. Lusk, L. E. Marcucci, S. C. Pieper, R. Schiavilla, M. Viviani *et al.*, *Phys. Rev. Lett.* **120**, 052503 (2018).
- [40] D. Lonardoni, J. Carlson, S. Gandolfi, J. E. Lynn, K. E. Schmidt, A. Schwenk, and X. B. Wang, *Phys. Rev. Lett.* **120**, 122502 (2018).
- [41] K. Schmidt and S. Fantoni, *Phys. Lett. B* **446**, 99 (1999).
- [42] J. Carlson, S. Gandolfi, F. Pederiva, S. Pieper, R. Schiavilla, K. Schmidt, and R. Wiringa, *Rev. Mod. Phys.* **87**, 1067 (2015).
- [43] B. S. Pudliner, V. R. Pandharipande, J. Carlson, S. C. Pieper, and R. B. Wiringa, *Phys. Rev. C* **56**, 1720 (1997).
- [44] D. Lacroix, A. Boulet, M. Grasso, and C.-J. Yang, *Phys. Rev. C* **95**, 054306 (2017).
- [45] A. Bulgac, M. Forbes, and P. Magierski, in *The Unitary Fermi Gas: From Monte Carlo to Density Functionals*, edited by W. Zwerger, The BCS-BEC Crossover and the Unitary Fermi Gas. Lecture Notes in Physics, Vol. 836 (Springer, Berlin, Heidelberg, 2012).
- [46] Michael McNeil Forbes, S. Gandolfi, and A. Gezerlis, *Phys. Rev. A* **86**, 053603 (2012).
- [47] W. G. Dawkins, J. Carlson, U. van Kolck, and A. Gezerlis, *Phys. Rev. Lett.* **124**, 143402 (2020).
- [48] Y. Castin and F. Werner, in *The BCS-BEC Crossover and the Unitary Fermi Gas*, Lecture Notes in Physics (Springer, Berlin, 2011), Vol. 127, p. 127.
- [49] M. Buraczynski and A. Gezerlis, *Phys. Rev. Lett.* **116**, 152501 (2016).
- [50] A. Sarsa, S. Fantoni, K. E. Schmidt, and F. Pederiva, *Phys. Rev. C* **68**, 024308 (2003).
- [51] S. Pilati, I. Zintchenko, and M. Troyer, *Phys. Rev. Lett.* **112**, 015301 (2014).
- [52] P. J. Reynolds, D. M. Ceperley, B. J. Alder, and W. A. Lester, *J. Chem. Phys.* **77**, 5593 (1982).
- [53] L. Riz, F. Pederiva, D. Lonardoni, and S. Gandolfi, *Particles* **3**, 706 (2020).
- [54] G. Palkanoglou and A. Gezerlis, *Universe* **7**, 24 (2021).
- [55] S. Athanassopoulos, E. Mavrommatis, K. Gernoth, and J. Clark, *Nucl. Phys. A* **743**, 222 (2004).
- [56] J. Carrasquilla and R. G. Melko, *Nat. Phys.* **13**, 431 (2017).
- [57] S. Raschka, [arXiv:1811.12808](https://arxiv.org/abs/1811.12808).
- [58] J. Carlson, S. Gandolfi, K. E. Schmidt, and S. Zhang, *Phys. Rev. A* **84**, 061602 (2011).
- [59] M. M. Forbes, S. Gandolfi, and A. Gezerlis, *Phys. Rev. Lett.* **106**, 235303 (2011).
- [60] S. Gandolfi, K. E. Schmidt, and J. Carlson, *Phys. Rev. A* **83**, 041601 (2011).
- [61] L. Luo and J. E. Thomas, *J. Low. Temp. Phys.* **154**, 1 (2008).
- [62] N. Navon, S. Nascimbene, F. Chevy, and C. Salomon, *Science* **328**, 729 (2010).
- [63] M. J. H. Ku, A. T. Sommer, L. W. Cheuk, and M. W. Zwierlein, *Science* **335**, 563 (2012).

- [64] G. Zürn, T. Lompe, A. N. Wenz, S. Jochim, P. S. Julienne, and J. M. Hutson, *Phys. Rev. Lett.* **110**, 135301 (2013).
- [65] P. Coleman, *Introduction to Many-Body Physics* (Cambridge University Press, Cambridge, 2015).
- [66] D. Davesne, A. Pastore, and J. Navarro, *Phys. Rev. C* **89**, 044302 (2014).
- [67] M. Buraczynski, N. Ismail, and A. Gezerlis, *Phys. Rev. Lett.* **122**, 152701 (2019).