Efficient learning of a one-dimensional density functional theory

M. Michael Denner¹, Mark H. Fischer¹, and Titus Neupert¹*Department of Physics, University of Zurich, Winterthurerstrasse 190, 8057 Zurich, Switzerland*

(Received 13 May 2020; revised 14 August 2020; accepted 17 August 2020; published 10 September 2020)

Density functional theory underlies the most successful and widely used numerical methods for electronic structure prediction of solids. However, it has the fundamental shortcoming that the universal density functional is unknown. In addition, the computational result—energy and charge density distribution of the ground state—is useful for electronic properties of solids mostly when reduced to a band structure interpretation based on the Kohn-Sham approach. Here, we demonstrate how machine learning algorithms can help to free density functional theory from these limitations. We study a theory of spinless fermions on a one-dimensional lattice. The density functional is implicitly represented by a neural network, which predicts, besides the ground-state energy and density distribution, density-density correlation functions. At no point do we require a band structure interpretation. The training data, obtained via exact diagonalization, feeds into a learning scheme inspired by active learning, which minimizes the computational costs for data generation. We show that the network results are of high quantitative accuracy and, despite learning on random potentials, capture both symmetry-breaking and topological phase transitions correctly.

DOI: [10.1103/PhysRevResearch.2.033388](https://doi.org/10.1103/PhysRevResearch.2.033388)

I. INTRODUCTION

Materials with strong electronic correlations host a variety of intriguing phenomena and quantum phases. Modeling and understanding these systems are among the greatest challenges in theoretical condensed matter physics. For quantitative and predictive results, numerical calculations are indispensable. The most widely and successfully used numerical approach to the electronic structure problem is based on density functional theory (DFT). In condensed matter physics, DFT is often linked to band structure calculations, while it is in principle much more powerful than that. The Hohenberg-Kohn theorems guarantee that a (potentially correlated) many-body ground state is uniquely determined by its energy and charge density distribution [1]. However, for practical implementations and a physical interpretation of calculated results, the Kohn-Sham ansatz is commonly used, producing the band structure of a different, noninteracting system with the same energy and density [2]. The implicit assumption is that this band structure captures the essential physics of the original system, at least if correlations are weak enough.

A critical shortcoming of DFT is that its eponymous functional is not known; instead, approximations on various levels of complexity are commonly employed [3]. It is important to emphasize that most of the functional is *universal*, representing the many-particle Schrödinger equation. The only nonuniversal input in a DFT calculation for a crystal is the

potential landscape within the unit cell induced from the ions and core electrons as well as the particle number, both of which do not affect the universal part of the functional.

The recent rise of machine-learning methods used to model physical systems has sparked hopes to use these methods for improving DFT calculations [4]. The approaches interject the DFT workflow at various stages, ranging from improving the Kohn-Sham scheme by representing the exchange-correlation functionals [5–10] or approximating the unknown energy functional and its derivatives [11–19]. More recent works bypass the Kohn-Sham-solution scheme by directly learning the mapping between material parameters and ground-state properties [20–31], or constructing the ground-state wave function from a corresponding density distribution [32]. Despite these advancements, previous approaches are either limited by complicated, nonscalable networks, suffer from inefficient training data generation or struggle in applications to the different physical phases of the used models.

In this work, we take an approach that follows three guiding principles: (i) implicit knowledge representation is the key strength of neural networks. Therefore we use a *neural network to implicitly represent the (minimized) DFT functional*. (ii) We aim at solving for phases of quantum matter beyond the band structure paradigm. To that end, we train the neural network to directly *output correlation functions* [32], which can be used to characterize phases and phase transitions. (iii) A balanced dataset is the key challenge as data acquisition—theoretical or experimental—is costly. In such settings, *active learning schemes* [33–37] offer better results with fewer training instances. In general, active learning describes an algorithm which can actively choose the data it wants to learn from during training. Here we employ a procedure inspired by active learning, to incorporate data from different system sizes in the a priori data generation. In particular, our approach uses costly data of larger

FIG. 1. (a) Schematic of the dense neural network used to learn the map from unit-cell potentials and filling to ground-state energy and density-density correlators. (b) Data generation inspired by active learning allows to check the energy deviations of different system sizes and continuing to larger systems only if necessary. (c) Exact versus predicted ground-state energies on the test data set of the actively learning network (ALN). The inset shows the absolute error on the test-energy values as a density histogram n . (d) Mean absolute error per correlator entry on the test data set for the ALN. (e) Exact versus predicted ground-state energies for the network trained on a single system size (PLN), with the inset showing the density n of absolute errors on the test data set. (f) Mean absolute error per correlator entry on the test data set for the PLN.

systems only in situations where large finite-size effects are detected.

Figures 1(a) and 1(b) summarize our model, neural network, and workflow. We choose to work with a one-dimensional lattice model of spinless fermions. The hopping and interaction terms of the model define our “universal Schrödinger equation” and are therefore left unaltered throughout the study. Input to the neural network is the problem-specific potential and particle number. Its output is the ground-state energy E_{GS} as well as the density-density correlation function. We start by introducing the employed learning scheme and demonstrate the quantitative accuracy of network predictions, after training it on random potentials. The active training shows superior performance compared to conventional training. We obtain mean squared errors of the energy of 3.08×10^{-4} in units of the hopping integral. Finally, we apply the trained network to a topological and a symmetry-breaking phase transition. Our results demonstrate a scalable architecture, able to capture interacting lattice models, with successful applications to structured phases.

II. MODEL

While density functional theory was originally formulated as a continuum theory, it has also been successfully applied to lattice models [38]. We consider a Hamiltonian for spinless fermions on a one-dimensional lattice with sites labeled by

$i = 1, \dots, L$ under periodic boundary conditions,

$$\hat{H} = -t \sum_i (\hat{c}_i^\dagger \hat{c}_{i+1} + \text{H.c.}) + U \sum_i \hat{n}_i \hat{n}_{i+1} + U' \sum_i \hat{n}_i \hat{n}_{i+2} + \sum_i V_i \hat{n}_i, \quad (1)$$

where $\hat{c}_i^\dagger$ and $\hat{c}_i$ are the fermion creation and annihilation operators on site i and $\hat{n}_i = \hat{c}_i^\dagger \hat{c}_i$ is the corresponding density operator. Nearest-neighbor hopping is parametrized by t , which will serve as the energy unit throughout. The particles are subject to a repulsive interaction on nearest- and next-nearest-neighbor sites, which we fix to $U = 1$ and $U' = 0.5$ so as to model a lattice analog of the Coulomb interaction. This parameter choice places the system in a metallic, but strongly correlated phase in absence of a potential V_i , even at half filling [39].

Motivated by the Hohenberg-Kohn theorems, we consider the kinetic term and the electron-electron interactions as universal, such that the external or ionic potential $\hat{V}_{\text{ext}} = \sum_i V_i \hat{n}_i$ together with the filling uniquely determine the ground state and all of its properties. We only consider potentials with periodicity of four sites. That is, the four values $V_i, i = 1, \dots, 4$, completely specify the Hamiltonian for any lattice size $L = 4N_{uc}$, with N_{uc} the number of unit cells and $V_i = V_{i+4}$ for all i . This four-site unit cell can be thought of as the discretized unit cell of a periodic crystal, while V_i is the ionic potential in this analogy. We restrict it to the range $V_i \in [-4, 4]$. We further

denote by the real number $0 < n_e < 4$ the particle filling per unit cell. We emphasize that despite imposing this periodicity to the potential, our approach is able to capture phases which spontaneously break the four-site translation symmetry.

III. LEARNING

The supervised-machine-learning algorithm we use bypasses the Kohn-Sham scheme by directly learning the map from the external parameters n_e and V_i to the corresponding ground-state energy and density-density correlators $\langle \hat{n}_i \hat{n}_j \rangle_{\text{GS}}$. The density-density correlators are calculated for two adjacent unit cells.¹ The chosen fully connected neural network [40] consists of four hidden layers which increase in size towards the output as depicted in Fig. 1(a) (see also Appendix A).

A central challenge in machine learning is unbiased and efficient data generation; one usually deals with limited computational or experimental resources. Here, we generate data by finite-size exact diagonalization (ED) of systems with randomly chosen n_e , V_i . In order to reduce finite size effects, these examples should naively be generated with as large systems as possible. However, the computational cost for data generation with ED grows exponentially with system size. For this reason, we employ a procedure inspired by *active learning*, performing costly large system ED, as depicted in Fig. 1(b), only if necessary. Using random values for n_e and V_i and starting with a comparison between $N_{\text{uc}} = 3, 4$, the scheme iteratively computes larger systems until the finite size deviation between ground-state energies lies below a priorly chosen threshold θ . Correspondingly, the fast computation of smaller systems is used as often as possible, while providing more accurate data in critical cases (see Appendix B). The samples are further augmented by applying translations within the unit cell and inversion, allowing the network to capture the symmetries of the universal part of the Hamiltonian.

We contrast the active learning approach outlined above with a passive learning scheme using training data generated for systems of fixed size $N_{\text{uc}} = 5$, with filling n_e and on-site potentials V_i chosen randomly. This system size is still solvable efficiently by ED, while sufficiently reducing finite size effects. The data are again symmetry-augmented. Both learning procedures were run with an equal time budget to ensure comparability.

A mean absolute error loss function is then optimized to obtain the weights and biases of the actively (ALN) and passively (PLN) learning neural network. The resulting performance is evaluated on unseen data, consisting of 20% of the full data set. Overfitting was avoided for both systems by suitable hyperparameter choices (see Appendix C). The absence of significant deviations in the energy correlation plot in Fig. 1(c) shows that the ALN performs well on random data, with an absolute error peaked at 1.2×10^{-2} . Similarly, Fig. 1(d) shows only small errors in the correlator prediction, with an overall mean absolute error of 1.8×10^{-3} . The PLN

performs worse in predicting energy values and correlators, as shown in the inset of Fig. 1(e) and in Fig. 1(f), with errors at least twice as large as in the case of the ALN. This comparison shows the advantage of intelligent data generation at fixed computational time budget.

Random potentials rarely represent a relevant physical scenario. We therefore present in the following how the randomly trained ALN outperforms the PLN for structured systems, where V_i obey further symmetries. Interestingly, including the smaller samples, which are suffering from finite-size effects, with a reduced sample weight in the ALN training leads to a more stable model prediction, in particular for physical systems (see Appendix B). We attribute this feature to a regularization of the model.

IV. LEARNABILITY OF OBSTRUCTED ATOMIC LIMITS

We consider the model introduced in Eq. (1) for a potential choice $(V_1, V_2, V_3, V_4) = (0, V, V, 0)$ at quarter filling ($n_e = 1$). The system is metallic for $V = 0$ which separates two distinct insulating phases for $V > 0$ and $V < 0$. For $V > 0$ ($V < 0$), Wannier functions are localized between the unit cells (in the middle of the unit cell). This corresponds to two topologically distinct (obstructed) atomic-limit insulators, with the intra-unit-cell hopping being effectively reduced (enhanced) compared to the inter-unit-cell hopping. The insulating nature of these phases can be shown by calculating the compressibility

$$\kappa = \frac{1}{n_e^2} \left(\frac{\partial^2 E_{\text{GS}}(n_e)}{\partial n_e^2} \right)^{-1}, \quad (2)$$

where n_e is the electron filling and $E_{\text{GS}}(n_e)$ the corresponding ground-state energy. Figure 2(a) shows that, as one approaches the critical metallic state around $V = 0$, κ increases rapidly. We emphasize that since κ is the second derivative of the energy, it is extremely susceptible to errors. Note, further, that the ED data show a strong even-odd effect in N_{uc} . We also calculated $\kappa(V)$ with the density matrix renormalization group algorithm (DMRG)² [41] for $N_{\text{uc}} = 28$. Compared with these exact results, the ALN produces a meaningful $\kappa(V)$ with a peak value $\kappa(V = 0)$ interpolating between the results of even and odd N_{uc} . On the contrary, the PLN is worse with a less pronounced and nonsymmetric peak. Even though only trained with at most $N_{\text{uc}} = 6$ data, we see that the ALN also resembles the $N_{\text{uc}} = 28$ DMRG result reasonably well.

The location of Wannier centers can be used to differentiate between the two phases. Defining $C = (\langle \hat{n}_2 \hat{n}_2 \rangle - \langle \hat{n}_2 \hat{n}_3 \rangle) - (\langle \hat{n}_4 \hat{n}_4 \rangle - \langle \hat{n}_4 \hat{n}_5 \rangle)$,³ the trivial atomic limit with localization in the unit cell is obtained for $C > 0$, the phase transition happens at $C = 0$ and the nontrivial atomic limit has $C < 0$.

²Calculations were performed using the TeNPy Library (version 0.5.0), for a finite but periodic lattice with maximum bond dimension $\chi_{\text{max}} = 2000$, max. truncation error $\epsilon = 10^{-8}$ and energy convergence criterion $\Delta E = 10^{-6}$.

³Note that we choose C to highlight that two neighboring sites are always occupied by just one electron. That distinguishes it from the band insulator case of filling $n_e = 2$ and a spontaneously translation symmetry broken phase at $n_e = 1$ (both $C = 0$).

¹In order to also use small systems ($N_{\text{uc}} = 4$) with periodic boundary conditions as part of the training data, at most two adjacent unit cells from the interior of the density-density correlator are a valid representative.

FIG. 2. Neural network results for a transition between different (obstructed) atomic limit insulators. (a) Compressibility κ for various potential strengths at quarter filling, as calculated from the actively and passively learned neural network, exact diagonalization (ED) and density matrix renormalization group (DMRG) of several system sizes. (b) Schematic depiction of the potential in the four-site unit cell: depending on the strength and sign, two obstructed atomic limits and a metallic phase can be realized. (c) Corresponding observable C as calculated from the 8×8 density-density correlator for the same numerical methods as used for the compressibility. The insets show the correlator as obtained from the ALN in the first (lower left) and second atomic limit (upper right) with the unit cell depicted in red.

Figure 2(c) highlights that both networks are able to capture C across the transition well, but the ALN results are markedly more accurate than the PLN results when compared with the ED and DMRG data. This supports the statement that our active learning scheme delivers quantitatively better results.

V. LEARNABILITY OF SPONTANEOUSLY SYMMETRY-BROKEN PHASES

Spontaneous breaking of translation symmetry can be triggered by introducing a potential of the form $(V_1, V_2, V_3, V_4) = (-V, V, -V, V)$ at quarter filling ($n_e = 1$). The symmetry-broken phase arises from the competition between the next-nearest-neighbor interaction U' and the increasing potential V . The four-site translational symmetry of Hamiltonian (1) is broken spontaneously at $V_c \approx 1.8$ (see Appendix D) by the two-site relations of the emerging degenerate ground-states. The metallic system shows a smooth dependence of E_{GS} on n_e around quarter filling [Fig. 3(a)]. With increasing

FIG. 3. Neural network results for a spontaneous symmetry-breaking phase. (a) Ground-state energy for $V = 0$ for several electron fillings n_e as calculated from the actively and passively learned neural network and ED. The inset displays the noninteracting band structure. (b) Ground-state energy as a function of the electron filling n_e in the symmetry broken phase ($V = 4$). The kink at $n_e = 1$ signals an interaction-induced incompressible phase. The inset reveals the band flattening of the noninteracting system. (c) Phase diagram, schematic of the potential, and density-density correlation functions. The latter are obtained from the ALN for $V = 0$ (left) and 4 (right). Off-diagonal terms, whose inequivalence signals the symmetry-breaking phase (right) are highlighted in red and green.

V , $E_{GS}(n_e)$ develops a kink at $n_e = 1$, signaling the emergence of the symmetry-broken charge-density wave (CDW) insulator [Fig. 3(b)]. Both neural networks represent the different phases very well, the deviations at small fillings are attributed to the limited amount of training samples in this limit.

Figure 3(c) shows that the correlation in the metallic phase is short-ranged and fast decaying, whereas the symmetry broken phase possesses a distinct order. The corresponding order parameter $\langle C_{SSB} \rangle = \frac{1}{N_{uc}} \langle \sum_i (-1)^i \hat{n}_{2i+1} \rangle$ is, however, zero, since the two degenerate ground states have opposite imbalance in electron density between the first and third sites of each unit cell. Instead, the order can be diagnosed by computing the square of the order parameter from the density-density correlation functions. This amounts to $\langle C_{SSB}^2 \rangle = \frac{2}{N_{uc}^2} \sum_{i \neq j} \langle \hat{n}_{4i+1} \hat{n}_{4j+1} \rangle - \frac{2}{N_{uc}^2} \sum_{i,j} \langle \hat{n}_{4i+1} \hat{n}_{4j+3} \rangle + C_0$, with an overall shift $C_0 = \frac{1}{N_{uc}^2} \langle \sum_i \hat{n}_{2i+1} \rangle$. Its nonzero value in the symmetry broken phase is implied by the inequivalence between the first (+) and second (−) term in the above expression,

highlighted in the density-density correlator in Fig. 3(c) by the red (+) and green (−) squares. This behavior is well captured by the ALN, producing quantitatively accurate correlations in both phases.

VI. CONCLUSION

We presented a supervised learning approach for lattice DFT, bypassing the Kohn-Sham solution scheme. Employing a procedure inspired by active learning allowed us to improve our results at fixed computational cost regarding data generation. Focussing on correlation functions on a subsystem and taking only the potential landscape in the unit cell and particle number as input results in a scalable architecture. Besides verification of our algorithm on unseen random potentials, we demonstrated that the trained networks reliably solve for different structured phases.

Looking ahead, it is highly desirable to construct similar implicit (neural network) representations of DFT for systems in continuous space and higher dimensions, in particular to attack the electronic structure problem in strongly correlated regimes. The main challenge is the generation of valid and balanced data sets, and the incorporation of data from various sources, including conventional DFT, Monte Carlo calculations, experiments, and future quantum simulation devices. Two concepts on which our study relies, (1) focus on correlation functions instead of quantum states and (2) the use of efficient learning, should prove useful in this future venture.

ACKNOWLEDGMENTS

We thank Giuseppe Carleo and Xi Dai for insightful discussions. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation program (ERC-StG-Neupert-757867-PARATOP).

APPENDIX A: NETWORK PARAMETERS

The supervised-machine-learning algorithm we propose in this paper uses dense neural networks of identical architecture, for both the active and passive learning scheme. The layers are connected by Softplus ($x = \ln(1 + e^x)$) activation functions, except for the output layer. The last layer takes into account that correlator values and ground-state energies have different ranges, by employing a linear activation function. Table I indicates the relevant parameters used in this paper.

APPENDIX B: TRAINING DATA

The choice of training examples is crucial in a machine learning setting, as data is precious. An intelligent data generation procedure is advantageous if computational time is finite, removing the necessity to always calculate as large systems as possible. Naively, the latter is the way to go in order to avoid finite size biases in the training examples. We contrast these two approaches as active and passive learning schemes.

The naive approach generates data with exact diagonalization of a five unit cell system, with electron fillings n_e

TABLE I. Relevant parameters used to create the neural networks in this paper.

Parameter	Value
Neurons layer 1	50
Activation 1	Softplus
Weight init. 1	lecun_uniform
Neurons layer 2	125
Activation 2	Softplus
Weight init. 2	lecun_uniform
Neurons layer 3	150
Activation 3	Softplus
Weight init. 3	lecun_uniform
Neurons layer 4	200
Activation 4	Softplus
Weight init. 4	lecun_uniform
Neurons layer 5	65
Activation 5	Linear
Weight init. 5	lecun_uniform
Optimizer	Adam
Batch size	100
Learning rate	0.001
Epochs	1500

and potentials V_i in the unit cell chosen randomly. The potentials are in the range $|V_i| \leq 4$, and the filling can assume values between $0.6 \leq n_e \leq 3.4$ electrons per unit cell. Data augmentation is then used to enlarge the dataset and allow the neural network to capture the underlying symmetries of the physical problem. This means applying translations $V_i \rightarrow V_{i+1}$ and inversion within the unit cell, where the former leads to a shift of the correlator data. The resulting dataset consists of 12,020 pairs of fillings and external potentials, mapped to the corresponding ground-state energies and density-density correlators. The split in training, validation and test sets is illustrated in Table II.

However, large system diagonalization is not always necessary to obtain accurate data. This is the basis for an active learning scheme, which performs the diagonalization of large systems only if strong finite size effects are detected. The input to this procedure is the random choice of electron fillings n_e and potentials V_i in the unit cell. The former assuming values between $0.5 \leq n_e \leq 3.5$ electrons per unit cell and the latter being chosen in the range $|V_i| \leq 4$. The training data for both networks was chosen as $|V_i| \leq 4$ in order to ensure that the transition to the symmetry broken phase lies within the trained potential range.

TABLE II. Data for the passively trained neural network.

Parameter	Value
Total number of samples	12020
Number of training samples	7212
Number of validation samples	2404
Number of test samples	2404
Number of samples from 20 site ED	12020
Sample weight	1

TABLE III. Data for the actively trained neural network.

Parameter	Value
Total number of samples	74 500
Number of training samples	44 700
Number of validation samples	14 900
Number of test samples	14900
Number of samples from 16 site ED	68 850
Number of samples from 20 site ED	5500
Number of samples from 24 site ED	150
Sample weight 16 site data	1 ^a
Sample weight 20 site data	2 ^b
Sample weight 24 site data	3

^aThis weight is increased to 3 if no 20 or 24 site system had to be calculated for this sample.

^bThis weight is increased to 3 if no 24 site system had to be calculated for this sample.

If the ground-state energy for a $N_{uc} = 3$ and $N_{uc} = 4$ unit cell system, calculated with exact diagonalization, deviates more than a priori chosen threshold θ , a system with one additional unit cell is being calculated. This procedure is repeated up to $N_{uc} = 6$ unit cells if necessary. This means that the neural network can query a larger system whenever the deviation of ground-state energies exceeds the chosen threshold θ , thereby reducing finite size effects. Naively one would remove the small inaccurate samples from the dataset, however we found that including them with a reduced sample weight yields similar results on random test data. However, we obtain better and more stable results on physical phases, meaning that the presence of the less accurate data points acts as a regularization of the model. As we strive for meaningful results on physical phases and not on random potentials, we consider this approach to be more promising. Data augmentation is again used not only to enlarge the dataset to 74.500 input-output data pairs (see Table III), but also to allow the network to capture the underlying symmetries.

The choice of the threshold θ is a crucial parameter of the active learning scheme. We therefore investigate the implications of different thresholds (see Table IV), and test the performance of the trained models on unseen examples. These examples are not part of the training set and were generated with random fillings n_e and potentials V_i in the range $|V_i| \leq 4$ for systems of $N_{uc} = 5, 6$.

The resulting mean absolute error is presented in Fig. 4, indicating that the error decreases with increasing θ . When considering random potentials, it is therefore advantageous to use training data generated with as large θ as possible,

TABLE IV. Number of samples for different choices of the threshold θ , calculated with exact diagonalization and an equal time budget for each parameter choice.

N_{uc}	$\theta = 0.0$	$\theta = 0.001$	$\theta = 0.0025$	$\theta = 0.005$	$\theta = 0.01$	$\theta = 1.5$
4	1350	4200	13 500	18 850	81 075	83 250
5	1335	575	1030	790	850	0
6	245	10	30	15	20	0

FIG. 4. Mean absolute error of neural network predictions for various datasets, trained on examples generated with different threshold θ as illustrated in Table IV. Random potentials were used to generate the evaluation data, with a $N_{uc} = 5$, $|V_i| \leq 4$ dataset consisting of 5600 samples, $N_{uc} = 6$, $|V_i| \leq 4$ containing 220 datapoints and 1205 $N_{uc} = 5$, $|V_i| \leq 1$ input - output pairs. The evaluation sets were not used in the training process of the neural networks.

resulting in the largest possible dataset. Consequently, this means larger systems with $N_{uc} = 5$ and 6 are never calculated.

The random data generation with potentials in the range $|V_i| \leq 4$ causes however mostly situations where the electrons are localized in the potential landscape, due to large potentials V_i . Finite size deviations are however negligible in such a case. This means that a good performance on this data does not necessarily correspond to a good elimination of finite size effects in a realistic physical potential. Figure 4 highlights this with the evaluation of a dataset with potentials $|V_i| \leq 1$, with the minimum of the obtained mean absolute error shifting towards smaller values of θ .

Consequently, a trade-off is needed between the total number of training samples and the number of samples with

TABLE V. Performance of the passively trained neural network. The last table section shows the performance on the ALN test dataset.

Parameter	Value
MAE on E_{GS}^{test}	2.8×10^{-2}
MSE on E_{GS}^{test}	1.5×10^{-3}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{test}$	3.6×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{test}$	5.9×10^{-5}
MAE on $E_{GS}^{validation}$	2.6×10^{-2}
MSE on $E_{GS}^{validation}$	1.3×10^{-3}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{validation}$	3.5×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{validation}$	4.9×10^{-5}
MAE on $E_{GS}^{training}$	2.5×10^{-3}
MSE on $E_{GS}^{training}$	1.2×10^{-3}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{training}$	3.2×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{training}$	4.2×10^{-5}
MAE on $E_{GS}^{ALNtest}$	3.7×10^{-2}
MSE on $E_{GS}^{ALNtest}$	2.7×10^{-3}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{ALNtest}$	4.5×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{ALNtest}$	7.0×10^{-5}

FIG. 5. (a) Finite size scaling plot of the critical potential V_c , obtained with a small potential δV to break the ground-state degeneracy. Density matrix renormalization group (DMRG) calculations of open boundary conditions are combined with the exact diagonalization (ED) of a 16 site system. (b) Schematic of the potential landscape in the unit cell. (c) Finite size scaling plot of the critical potential V_c , obtained with an additional half unit cell and electron to break the ground-state degeneracy. Density matrix renormalization group (DMRG) calculations of open boundary conditions are combined with the exact diagonalization (ED) of an 18 site system. (d) Schematic of the potential landscape with additional two lattice sites.

$N_{uc} = 5, 6$. The best compromise is reached for $\theta = 0.0025$, with a large enough size of the dataset, while containing significant information about larger systems. This parameter was therefore used to generate the final dataset used to train the actively learning neural network in this paper.

TABLE VI. Performance of the actively trained neural network. The last table section shows the performance on the PLN test dataset.

Parameter	Value
MAE on E_{GS}^{test}	1.2×10^{-2}
MSE on E_{GS}^{test}	2.4×10^{-4}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{test}$	1.8×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{test}$	1.8×10^{-5}
MAE on $E_{GS}^{validation}$	1.2×10^{-2}
MSE on $E_{GS}^{validation}$	2.3×10^{-4}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{validation}$	1.8×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{validation}$	1.5×10^{-5}
MAE on $E_{GS}^{training}$	1.2×10^{-2}
MSE on $E_{GS}^{training}$	2.3×10^{-4}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{training}$	1.8×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{training}$	1.2×10^{-5}
MAE on $E_{GS}^{PLNtest}$	1.8×10^{-2}
MSE on $E_{GS}^{PLNtest}$	5.2×10^{-4}
MAE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{PLNtest}$	2.9×10^{-3}
MSE on $\langle \hat{n}_i \hat{n}_j \rangle_{GS}^{PLNtest}$	3.6×10^{-5}

APPENDIX C: TRAINING PERFORMANCE

The actively and passively constructed datasets are split into training, validation and test sets. Achieving consistent performance on the optimized and unseen data indicates that the network has not been overfitted. The corresponding results for both training schemes are presented in Tables V and VI, highlighting that the intended mapping from electron fillings and potentials to ground-state energy and density-density correlator is well captured.

APPENDIX D: DETERMINATION OF THE POINT OF SPONTANEOUS SYMMETRY-BREAKING

The four-site translational symmetry of the considered extended Hubbard model can be spontaneously broken by the introduction of an external potential $\hat{V}_{ext} = \sum_{i=1}^4 V_i \hat{n}_i$. Choosing a potential of the form $V_1 = -V_2 = V_3 = -V_4 = -V$ causes a competition between next-nearest neighbor repulsion U' and the potential V at quarter filling. As a result, a symmetry-breaking phase with two degenerate ground-states emerges, corresponding to ordering the electrons either to site one or to site three in each unit cell.

We consider two approaches in order to identify the potential strength at which the spontaneous symmetry-breaking occurs. For each of these, we investigate a finite size scaling plot, to derive the extrapolated point of the phase transition in the thermodynamic limit.

As stated above, the symmetry broken phase possesses a distinct order in the density-density correlator, occupying only

sites with negative potential. By adding a small potential on one site, e.g., $V_1 = -V - \delta V$, $\delta V \ll V$, one of the degenerate ground-states is favoured. Consequently, the occupation $\langle \hat{n}_i \rangle_{\text{GS}}$ will concentrate on the first lattice site of each unit cell. We therefore investigate the difference in the occupation of site one and three, $C = \langle \hat{n}_1 \rangle_{\text{GS}} - \langle \hat{n}_3 \rangle_{\text{GS}}$. In the metallic phase for a vanishing potential $V = 0$, all sites are equally occupied, despite the small offset δV . This changes around the point of the phase transition, where explicitly one of the symmetry-breaking ground states is selected. The position of this jump in C can then be studied with various system sizes and extrapolated to the thermodynamic limit. Since exact diagonalization already reaches computational boundaries for quite small systems, we additionally employ the density-matrix renormalization group algorithm (DMRG)⁴ [41]. In order to ensure convergence also for large systems, open boundary conditions

are considered. The influence of the boundary can however be neglected when considering density-density correlations in the middle of the system. Figure 5(a) indicates the transition point as extrapolated from several system sizes, leading to a potential of $V_c = 1.744$.

Additionally, one can study the emergence of the spontaneously symmetry-breaking phase in the energy spectrum. The restriction to systems with open boundary conditions for DMRG calculations leads to a degeneracy of ground states of the order of the number of unit cells, compared to two states in the case of periodic boundaries. We therefore enlarge the system by two lattice sites and add one additional electron for the entire system, breaking the ground-state degeneracy. The transition to the symmetry broken phase can consequently be probed by calculating the gap $E_{\text{GS}} - E_1$ for several system sizes. The critical potential V_c corresponds to the point of gap opening. The corresponding results are presented in Fig. 5(c), with an extrapolated transition point at $V_c = 1.864$. Taking the mean between both results gives a critical potential $V_c \approx 1.8$. Correspondingly, choosing the range of training potentials as $V_i \in [-4, 4]$ secures that the neural networks can capture the phase transition.

⁴Calculations were performed using the TENPY Library (version 0.5.0), for a finite lattice with maximum bond dimension $\chi_{\text{max}} = 1000$, max. truncation error $\epsilon = 10^{-10}$ and energy convergence criterion $\Delta E = 10^{-6}$.

-
- [1] P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964).
  - [2] W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965).
  - [3] P. Mori-Sánchez, A. J. Cohen, and W. Yang, *Phys. Rev. Lett.* **100**, 146401 (2008).
  - [4] G. R. Schleder, A. C. M. Padilha, C. M. Acosta, M. Costa, and A. Fazzio, *J. Phys.: Mater.* **2**, 032001 (2019).
  - [5] K. T. Lundgaard, J. Wellendorff, J. Voss, K. W. Jacobsen, and T. Bligaard, *Phys. Rev. B* **93**, 235162 (2016).
  - [6] B. Kolb, L. C. Lentz, and A. M. Kolpak, *Sci. Rep.* **7**, 1192 (2017).
  - [7] Q. Liu, J. Wang, P. Du, L. Hu, X. Zheng, and G. Chen, *J. Phys. Chem. A* **121**, 7273 (2017).
  - [8] R. Nagai, R. Akashi, S. Sasaki, and S. Tsuneyuki, *J. Chem. Phys.* **148**, 241737 (2018).
  - [9] S. Dick and M. Fernandez-Serra, *J. Chem. Phys.* **151**, 144102 (2019).
  - [10] J. Schmidt, C. L. Benavides-Riveros, and M. A. L. Marques, *J. Phys. Chem. Lett.* **10**, 6425 (2019).
  - [11] J. C. Snyder, M. Rupp, K. Hansen, K.-R. Müller, and K. Burke, *Phys. Rev. Lett.* **108**, 253002 (2012).
  - [12] J. C. Snyder, M. Rupp, K. Hansen, L. Blooston, K.-R. Müller, and K. Burke, *J. Chem. Phys.* **139**, 224104 (2013).
  - [13] L. Li, J. C. Snyder, I. M. Pelaschier, J. Huang, U.-N. Niranjan, P. Duncan, M. Rupp, K.-R. Müller, and K. Burke, *Int. J. Quantum Chem.* **116**, 819 (2016).
  - [14] L. Li, T. E. Baker, S. R. White, and K. Burke, *Phys. Rev. B* **94**, 245129 (2016).
  - [15] K. Yao and J. Parkhill, *J. Chem. Theory Comput.* **12**, 1139 (2016).
  - [16] J. Seino, R. Kageyama, M. Fujinami, Y. Ikabata, and H. Nakai, *J. Chem. Phys.* **148**, 241705 (2018).
  - [17] J. Nelson, R. Tiwari, and S. Sanvito, *Phys. Rev. B* **99**, 075132 (2019).
  - [18] P. Golub and S. Manzhos, *Phys. Chem. Chem. Phys.* **21**, 378 (2019).
  - [19] T. Nudējima, Y. Ikabata, J. Seino, T. Yoshikawa, and H. Nakai, *J. Chem. Phys.* **151**, 024104 (2019).
  - [20] K. Hansen, G. Montavon, F. Biegler, S. Fazli, M. Rupp, M. Scheffler, O. A. von Lilienfeld, A. Tkatchenko, and K.-R. Müller, *J. Chem. Theory Comput.* **9**, 3404 (2013).
  - [21] K. T. Schütt, H. Glawe, F. Brockherde, A. Sanna, K. R. Müller, and E. K. U. Gross, *Phys. Rev. B* **89**, 205118 (2014).
  - [22] K. Hansen, F. Biegler, R. Ramakrishnan, W. Pronobis, O. A. von Lilienfeld, K.-R. Müller, and A. Tkatchenko, *J. Phys. Chem. Lett.* **6**, 2326 (2015).
  - [23] K. T. Schütt, F. Arbabzadah, S. Chmiela, K. R. Müller, and A. Tkatchenko, *Nat. Commun.* **8**, 13890 (2017).
  - [24] F. Brockherde, L. Vogt, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, *Nat. Commun.* **8**, 872 (2017).
  - [25] M. Bogojeski, F. Brockherde, L. Vogt-Maranto, L. Li, M. E. Tuckerman, K. Burke, and K.-R. Müller, *arXiv:1811.06255*.
  - [26] K. Ryczko, K. Mills, I. Luchak, C. Homenick, and I. Tamblyn, *Comput. Mater. Sci.* **149**, 134 (2018).
  - [27] E. Schmidt, A. T. Fowler, J. A. Elliott, and P. D. Bristowe, *Comput. Mater. Sci.* **149**, 250 (2018).
  - [28] S. Pilati and P. Pieri, *Sci. Rep.* **9**, 5613 (2019).
  - [29] C. A. Custodio, E. R. Filletti, and V. V. França, *Sci. Rep.* **9**, 1886 (2019).
  - [30] K. Ryczko, D. A. Strubbe, and I. Tamblyn, *Phys. Rev. A* **100**, 022512 (2019).
  - [31] L. Zepeda-Núñez, Y. Chen, J. Zhang, W. Jia, L. Zhang, and L. Lin, *arXiv:1912.00775*.
  - [32] J. R. Moreno, G. Carleo, and A. Georges, *Phys. Rev. Lett.* **125**, 076402 (2020).

- [33] B. Settles, Active learning literature survey, Computer Sciences Technical Report 1648, University of Wisconsin–Madison, 2009.
- [34] K. Gubaev, E. V. Podryabinkin, G. L. Hart, and A. V. Shapeev, *Comput. Mater. Sci.* **156**, 148 (2019).
- [35] G. Sivaraman, A. N. Krishnamoorthy, M. Baur, C. Holm, M. Stan, G. Csányi, C. Benmore, and Á. Vázquez-Mayagoitia, *npj Comput. Mater.* **6**, 104 (2020).
- [36] J. Yao, Y. Wu, J. Koo, B. Yan, and H. Zhai, *Phys. Rev. Res.* **2**, 013287 (2020).
- [37] G. H. Teichert, A. R. Natarajan, A. Van der Ven, and K. Garikipati, *Comput. Methods Appl. Mech. Eng.* **371**, 113281 (2020).
- [38] K. Schönhammer, O. Gunnarsson, and R. M. Noack, *Phys. Rev. B* **52**, 2504 (1995).
- [39] L. Markhof, B. Sbierski, V. Meden, and C. Karrasch, *Phys. Rev. B* **97**, 235126 (2018).
- [40] F. Chollet *et al.*, Keras, <https://keras.io> (2015).
- [41] J. Hauschild and F. Pollmann, *SciPost Phys. Lect. Notes* **5** (2018), code available from <https://github.com/tenpy/tenpy>.