SCIENTIFIC REP{O}RTS

OPEN

Machine learning quantum phases of matter beyond the fermion sign problem

Peter Broecker[1], Juan Carrasquilla[2], Roger G. Melko[2,3] & Simon Trebst [1]

State-of-the-art machine learning techniques promise to become a powerful tool in statistical mechanics via their capacity to distinguish different phases of matter in an automated way. Here we demonstrate that convolutional neural networks (CNN) can be optimized for quantum many-fermion systems such that they correctly identify and locate quantum phase transitions in such systems. Using auxiliary-field quantum Monte Carlo (QMC) simulations to sample the many-fermion system, we show that the Green's function holds sufficient information to allow for the distinction of different fermionic phases via a CNN. We demonstrate that this QMC + machine learning approach works even for systems exhibiting a severe fermion sign problem where conventional approaches to extract information from the Green's function, e.g. in the form of equal-time correlation functions, fail.

In quantum statistical physics, the sign problem refers to the generic inability of quantum Monte Carlo (QMC) approaches to tackle fermionic systems with the same unparalleled efficiency it exhibits for unfrustrated bosonic systems. At its most basic level, it tracks back to the expansion of the partition function of a quantum mechanical system in terms of $d + 1$ dimensional classical configurations that have both positive and negative (or complex) statistical weights, thus invalidating their usual interpretation as a probability distribution[1,2]. In some specific cases, canonical transformations or basis rotations are known that completely eliminate the negative weights[3–5], resulting in sign-problem free models, sometimes called "stoquastic"[6] or "designer"[7] Hamiltonians. However, the lack of a general systematic procedure for such transformations[8] preclude many, if not most, quantum Hamiltonians from being simulated with unbiased QMC methods. This includes one of the most fundamental problems in statistical physics – the many-electron system, which is known to give rise to some of the most intriguing collective phenomena such as the formation of high-temperature superconductors[9], non-Fermi liquids[10,11], or Mott insulators with fractionalized excitations[12].

When tackling sign-problematic Hamiltonians with QMC approaches a common procedure[13] consists of two steps: (1) taking the absolute value of the configuration weight, thereby allowing interpretation as a probability amenable to sampling; and (2) precisely compensating for this by weighing observables (such as two-point correlation functions) with the excluded sign. While this procedure allows, in principle, for an unbiased evaluation of observables, it introduces changes into the sampling scheme in two distinct ways. First, the exclusion of the sign in step (1) affects the region of configuration space that is effectively sampled. To what extent this modified sampling imposes severe restrictions or rather subtle constraints very much depends on the actual QMC flavor, such as world-line versus auxiliary-field approaches. Second, this modified sampling necessitates the sign reweighing of step (2) in any proper statistical analysis. It is, however, precisely this step where the sign problem ultimately manifests itself in a statistical variance of estimators that grows *exponentially* in system size and inverse temperature.

In this paper, we examine an approach by which these two steps in the sampling procedure of sign-problematic QMC can be separated in the context of the many-fermion problem. To do this, we replace step (2), the calculation of thermodynamic observables, with supervised machine learning on configuration data produced in step (1). Neural networks have recently been demonstrated capable of discriminating between classical phases of matter, through direct training on Monte Carlo configurations[14,15]. In this paper, we employ auxiliary-field QMC techniques to sample statistical instances of the wavefunction of a fermionic system. We then train a convolutional neural network (CNN) to discriminate between two fermionic phases, which are known ground states for

[1]Institute for Theoretical Physics, University of Cologne, 50937, Cologne, Germany. [2]Perimeter Institute for Theoretical Physics, Waterloo, Ontario, N2L 2Y5, Canada. [3]Department of Physics and Astronomy, University of Waterloo, Ontario, N2L 3G1, Canada. Correspondence and requests for materials should be addressed to S.T. (email: trebst@thp.uni-koeln.de)

certain parameters of a fermionic quantum lattice model, directly with QMC samples of the Green's function. Once trained, the CNN can provide a prediction, for instance, of the parametric location of the phase transition between the two phases, which we demonstrate for a number of Hubbard-type quantum lattice models with competing itinerant and charge-ordered phases. Importantly, this robust prediction of quantum critical points appears to work even for systems where the Monte Carlo sampling of conventional observables is plagued by a severe sign problem. Such a machine learning approach to the QMC sampling of many-fermion systems thus allows one to determine whether crucial information about the ground state of the many-fermion system is truly lost in the sampling procedure, or whether it can be retrieved in physical entities beyond statistical estimators, enabling a supervised learning of phases despite the presence of the sign problem.

Circumventing the Fermion-Sign Problem

To begin, consider a d-dimensional fermionic quantum system, which can be generically written in terms of a classical statistical mechanics problem defined on a phase space with configurations C in $d+1$ dimensions. The partition function of the quantum system can thereby be expressed as a sum of statistical weights over classical configurations, i.e. $Z = \sum_C W_C$. Unlike classical systems, for quantum Hamiltonians the weights W_C can be both positive and negative (or even complex), which invalidates the usual Monte Carlo interpretation of W_C/Z as a probability distribution. In principle, a stochastic interpretation can be salvaged by considering a modified statistical ensemble with probability distribution $P_C \propto |W_C|$ and concomitantly moving the sign of W_C to the observable

$$\langle O \rangle = \frac{\sum_C O_C \cdot W_C}{\sum_C W_C} = \frac{\sum_C O_C \cdot \text{sign}(W_C) \cdot |W_C|}{\sum_C \text{sign}(W_C) \cdot |W_C|}$$
$$= \frac{\langle \text{sign} \cdot O \rangle_{|W|}}{\langle \text{sign} \rangle_{|W|}}. \tag{1}$$

This procedure, although formally exact, introduces the QMC sign problem as a manifestation of the "small numbers problem", where the numerator and denominator in the last expression both approach zero exponentially in system size N and inverse temperature β[1, 2], i.e. we have

$$\langle \text{sign} \rangle_{|W|} = \exp(-\beta N \Delta f), \tag{2}$$

where Δf is the difference in the free energy densities of the original fermionic system and the one with absolute weights. Thus resolving the ratio in Eq. (1) within the statistical noise inherent to all Monte Carlo simulations becomes exponentially hard. The advantage of importance sampling, which often translates into polynomial scaling, is lost.

In this work, instead of attempting to obtain exact expectation values of physical observables, or attempting to find a basis where the weights W_C are always non-negative or that ameliorates the calculation of $\langle \text{sign} \rangle_{|W|}$, we introduce a basis-dependent "state function" F_C whose goal is to associate configurations C with the most likely phase of matter they belong to for a given Hamiltonian. More precisely, we assume that there exists a function F_C such that its expectation value in the modified ensemble of absolute weights

$$\langle F \rangle_{|W|} = \frac{\sum_C F_C \cdot |W_C|}{\sum_C |W_C|} \tag{3}$$

is 1 when the system is deep in phase A and 0 when the system is deep in the neighboring phase B. Around the critical point separating phase A from B, $\langle F \rangle_{|W|}$ crosses over from one to zero. The value $\langle F \rangle_{|W|} = 1/2$ indicates that the function can not make a distinction between phases A and B, and therefore assigns equal probability to both phases. We therefore interpret this value as locating the position of the transition separating the two phases in parameter space. In practice, we use a deep CNN to approximate the state function F, which is trained on "image" representations of configurations C sampled from the modified ensemble $|W_C|/\sum_C |W_C|$ in the two different phases A and B. We explore several choices for this image representation including color-conversions of the auxiliary field encountered in determinental Monte Carlo approaches, the Green's function as well as the Green's function multiplied by the sign. If the above procedure indeed allows the crafting of such a state function F, then one has found a path to a *sign-problem avoiding* discrimination of the two phases and their phase transitions through the evaluation of $\langle F \rangle_{|W|}$.

Convolutional Neural Networks

Artificial neural networks have for some time been identified as the key ingredient of powerful pattern recognition and machine learning algorithms[16, 17]. Very recently, neural networks and other machine learning algorithms have been brought to the realm of quantum and classical statistical physics[18–28]. On a conceptual level, parallels between deep learning and renormalization group techniques have been explored[29, 30], while on a more practical level machine learning algorithms have been applied to model potential energy surfaces[31], relaxation in glassy liquids[32] or the identification of phase transitions in classical many-body systems[14, 15]. Boltzmann machines, as well as their quantum extensions[33], have been applied to statistical mechanics models[34] and quantum systems[35]. In addition, new supervised learning algorithms inspired by tensor-network representations of quantum states have been recently proposed[36].

In machine learning, the goal of artificial neural networks is to learn to recognize patterns in a (typically high dimensional) data set. CNNs, in particular, are nonlinear functions which are optimized (in an initial "training" step) such that the resulting function F allows for the extraction of patterns (or "features") present in the data.
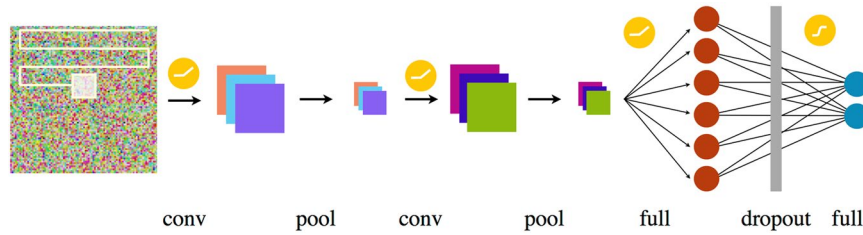
Figure 1. Schematic illustration of the neural network used in this work. A combination of convolutional (conv) and max pooling layers (pool) is first used to study the image, before the data is further analyzed by two fully connected neural networks separated by a dropout layer. The convolutional and the first fully connected layer are activated using rectified linear functions, while the final layer is activated by a softmax function.

Here we take this approach to construct a function F, represented as a deep CNN, that allows the classification of many-fermion phases as outlined in the previous section. Our choice of employing a deep CNN is rooted in the above observation that the configurations generated from a quantum Monte Carlo algorithm can be often interpreted as "images". As we explain below in more detail, our analysis can be regarded as an image classification problem – an extremely successful application of CNNs.

The architecture of the CNN we use is depicted schematically in Fig. 1 with a more detailed technical discussion of the individual components presented in the Methods section. We feed the CNN with Monte Carlo configurations (illustrated on the left), which, processed through the network, provide a two-component softmax output layer (on the right). The two components of this function, which by construction always add up to one, can be interpreted as the probabilities that a given configuration belongs to the two different phases under consideration and can thus be used for classification. In the initial training step, we optimize the CNN on a set of (typically) 2×8192 representative configurations sampled deep in the two fermionic phases. The question of which fundamental features, contained in the Monte Carlo configurations, are used in the resulting function F to characterize the phases under consideration, is automatically discovered during the training procedure (and beyond our direct influence).

Machine learning fermionic quantum phases

We apply this QMC + machine learning framework to a family of Hubbard-like fermion models where the competition between kinetic and potential terms gives rise to a phase transition between an itinerant metallic phase and a charge-ordered Mott insulating phase. As a first example we consider a system of spinful fermions on the honeycomb lattice subject to the Hamiltonian

$$H = K + V = -t \sum_{\langle i,j \rangle, \sigma} c_{i,\sigma}^\dagger c_{j,\sigma} + U \sum_i n_{\uparrow,i} n_{\downarrow,i},$$

(4)

with a kinetic term K and a potential term V. At zero temperature and half-filling, this system is well known to undergo a quantum phase transition from a Dirac semi-metal to an insulator with antiferromagnetic spin-density wave (SDW) order at $U_c/t \approx 3.85$[37]. For convenience, we will set $t = 1$ in the following.

To sample configurations for different values of the tuning parameter U we employ determinantal quantum Monte Carlo (DQMC) in its projective zero-temperature formulation. In this scheme, a carefully chosen test wave function $|\psi_T\rangle$ is projected onto the actual ground state function $|\psi\rangle$

$$|\psi\rangle = e^{-\theta H} |\psi_T\rangle.$$

(5)

To compute this projection, we first apply a Trotter decomposition to discretize the projection time θ into $N_\tau = \theta/\Delta\tau$ time steps and separate the kinetic and potential terms

$$e^{-\theta H} = \prod_{n=1}^{N_\tau} e^{-\Delta\tau K} e^{-\Delta\tau V} \equiv B(\theta).$$

(6)

The quartic interaction term is then decomposed by applying a Hubbard-Stratonovich (HS) transformation on each on-site interaction V_i and on each time slice τ

$$e^{-\Delta\tau V_i} = \frac{1}{2} \sum_{s=\pm 1} \prod_{\sigma=\uparrow,\downarrow} e^{-V_i(s,\tau,\sigma)},$$

(7)

introducing one auxiliary variable $s = \pm 1$ per site and separating the two spin species σ. The entirety of the auxiliary variables makes up the Hubbard Stratonovich field and will be denoted as **s** in the following. The probability for choosing a configuration is given by

$$p(\mathbf{s}, \mathbf{s}') = \frac{\langle \psi(\mathbf{s}) | \psi(\mathbf{s}') \rangle}{\sum_{\mathbf{ss}} \langle \psi(\mathbf{s}) | \psi(\mathbf{s}') \rangle},$$
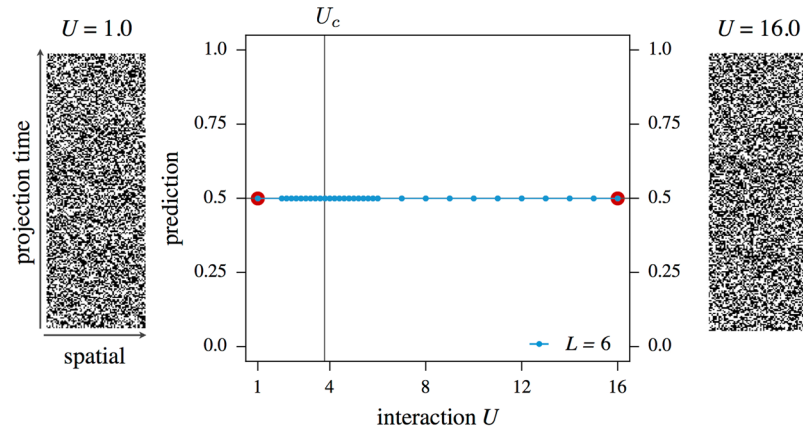
(8)

Figure 2. Results from training the neural network on Hubbard-Stratonovich field configurations of a spinful Hubbard model on a $2 \cdot 6 \times 6$ lattice with on-site interaction U. Reference points for training were $U = 1.0$ and $U = 16.0$, marked by red dots in the figure. Despite intensive training, the network depicted in Fig. 1 is unable to distinguish the auxiliary field configurations of the two reference points and as a consequence can not be used to discriminate between the two phases.

where $\mathbf{s}$ and $\mathbf{s}'$ denote the Hubbard-Stratonovich fields associated with the projection of the wavefunction used as bra and ket, respectively. The weight of the configuration $\langle\psi(\mathbf{s})|\psi(\mathbf{s}')\rangle$ evaluates to a determinant

$$\langle\psi(\mathbf{s})|\psi(\mathbf{s}')\rangle = \det[P^\dagger B(\theta, \mathbf{s}')B(\theta, \mathbf{s}')P] \qquad (9)$$

where P is the matrix representation of the test wave function $|\psi_T\rangle$. For auxiliary field approaches the modified statistical ensemble of absolute weights implies that the sign of the fermionic determinant will be ignored – importantly, such a modified ensemble *retains* the fermionic exchange statistics, but becomes insensitive to the parity of the total number of fermionic exchanges for a given configuration (which is precisely what is reflected in the sign of the determinant). This should be contrasted to world-line QMC approaches where the modified ensemble weighted by $|W_C|$ would not preserve any fermionic exchange statistics at all, but effectively sample a *bosonic* system.

In order to implement our machine learning approach, we begin by choosing the classical configuration space C over which the expectation values in Eqs (1) and (3) are taken. An obvious candidate is the auxiliary field s. This approach is illustrated in Fig. 2, where the CNN has been trained at parameters $U = 1$ and $U = 16$, i.e. deep within the Dirac semi-metal and the antiferromagnetic SDW phase, respectively. The side panels show representative reference configurations of the auxiliary field at each of these training parameters. Interestingly, the configurations displayed in Fig. 2 show no discernible difference between the two auxiliary field configurations, apparent to the human eye. Indeed, we find that optimizing the CNN of Fig. 1 to extract information directly from these auxiliary field configurations does not yield a function F that allows one to distinguish between the two phases. This apparent inability is possibly rooted in the particular choice of the employed Hubbard-Stratonovich transformation, which preserves SU(2) spin symmetry by decoupling in the charge channel. In the supplementary material we discuss an alternative Hubbard-Stratonovich transformation by decoupling in the spin channel (which does not preserve the SU(2) spin symmetry), which for the phase transition at hand also does not lead to satisfactory results. While it is well known on general grounds[38, 39] that the auxiliary field *can* reflect physical correlations (and as such should be amenable to the applied pattern recognition technique[40]) if the Hubbard-Stratonovich transformation is performed in the right channel, our goal here is to identify a somewhat more general approach that relies on more generic physical quantities.

To alleviate this difficulty, we instead consider the Green's function $G(i, j) = \langle c_i c_j^\dagger \rangle$ as input for our machine learning approach. The Green's function is an essential quantity in statistical physics, which allows e.g. for the calculation of equal-time correlation functions, and while it can easily be calculated from a given auxiliary field configuration it is not sensitive to the specifics of the Hubbard-Stratonovich transformation. Instead of the bare auxiliary fields, we thus train the CNN on the unprocessed complex valued Green's matrices $G_s(i, j) = \langle c_i c_j^\dagger \rangle_s$ calculated for a given auxiliary field configuration s. For the training, we used 2×8192 (2×4096 for $L = 15$) samples of the Green's function. This modified approach gives a striking improvement in results, as illustrated in Fig. 3. The side panels now show representative examples of the Green's matrices $G_s(i, j)$ for the two coupling parameters well inside the two respective fermionic phases. For the purpose of visualization, we convert the complex-valued entries of the Green's matrices to a polar representation which are then interpreted as HSV colors and finally converted to RGB for illustration[41]. Contrary to the visual inspection of the auxiliary field configuration in Fig. 2, the image-converted Green's function exhibits a clearly visible distinction for the two phases. Indeed the CNN trained and applied to the image-converted Green's function now succeeds in discriminating the two phases by producing a function F that indicates a phase transition around a value of the interaction $U \approx 4.1 \pm 0.1$. For a given finite system size L, we identify the location of the phase transition with the parameter U for which the averaged state function F is 1/2, i.e. the parameter for which the CNN cannot make any distinction between the
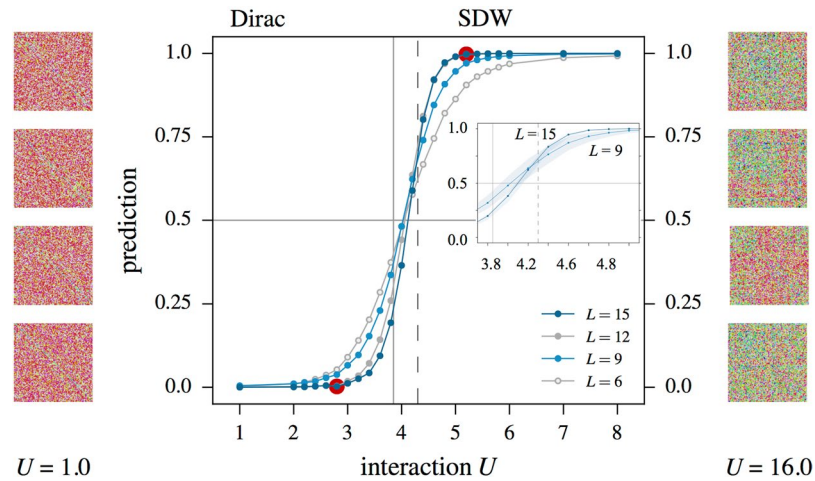
Figure 3. Machine learning of the phase transition from a semi-metal to an antiferromagnetic insulator in the spinful Hubbard model (4) on a honeycomb lattice using the Green's function approach (see main text). Visualized in the side panels are representative samples of the Green's function (calculated from the auxiliary field) for a $2 \cdot 9 \times 9$ system in the two respective phases. The complex entries of these matrices are color-converted by interpreting their absolute value as the hue of the color while their angle is chosen as the saturation (HSV coloring scheme[41]). The main panel shows the output of the discriminating function F obtained from a CNN trained for parameters in the two fermionic phases (indicated by the red dots). Data for different system sizes $2 \cdot L \times L$ are shown where the colors were selected to highlight an apparent even-odd effect in the linear system size. The vertical solid line indicates the position of the phase transition in the thermodynamic limit[37], while the dashed line marks the position at which the antiferromagnetic order breaks down[42] for the finite system sizes of the current study.

two phases and therefore assigns equal probability to both phases. These estimates for the location of the phase transition and their finite-size trends are in good agreement with the critical value of $U_c(L = 15) \approx 4.3$ obtained from Monte Carlo simulations for similar system sizes[42] and slightly above the critical value $U_c(L \rightarrow \infty) \approx 3.85$ of the thermodynamic limit[37].

Sign-problematic many-fermion systems

We now turn to many-fermion systems that exhibit a sign problem in the conventional QMC + statistical analysis approach, and ask to what extent the QMC + machine learning framework is sensitive to this sign problem. Simple example systems of this sort are *spinless* fermion models, which typically exhibit a severe sign problem in the conventional complex fermion basis (as we will illustrate below). We first consider a half-filled honeycomb system subject to the Hamiltonian,

$$H = -t\sum_{\langle i,j \rangle} c_i^\dagger c_j + V\sum_{\langle i,j \rangle} n_i n_j. \tag{10}$$

The competition between the kinetic term (which we again set to $t = 1$) and a repulsive nearest neighbor interaction V drives the system through a quantum phase transition[43] separating a semi-metallic state for $V < V_c$ from a charge density wave (CDW) state for $V > V_c$. Interestingly, this model can be made to be sign-problem free through a basis transformation to a Majorana fermion basis[5] or by using a continuous time quantum Monte Carlo flavor[44], which allows for a precise estimation of the critical repulsion $V_c \approx 1.36$ directly from QMC observables[44–48]. For the purpose of this paper, we will not perform this transformation, but rather sample the model in its sign-problematic formulation in the conventional complex fermion basis. The average sign, which here enters as a complex number, is illustrated for a range of couplings in the lower panel of Fig. 4. With a vanishingly small expectation value of the sign, we indeed encounter a severe sign problem.

Analogous to our procedure for the sign-problem free case of spinful fermions, we first train the CNN on representative samples of the Green's function for parameters deep within the two phases. To do so, we generate 8192 (4096 for $L = 15$) labeled samples for $V = 0.1$ (semi-metal) and $V = 2.5$ (CDW) from DQMC simulations using the modified statistical ensemble of absolute weights $|W_C|$ and train the CNN on these labeled instances. The trained CNN we then feed with *unlabeled* configurations from several different interaction values $0.1 < V < 2.5$ and ask the neural network to predict to which phase a particular configuration belongs.

At this point, a decision has to be made about how to provide information about the sign of each configuration to the CNN. We explore two options. First, we multiply each Green's matrix $G_s(i, j)$ with the sign (in general a complex phase) associated with the underlying configuration, i.e. sign(W_C), for a given Markov chain. Second, we ignore the sign altogether, and feed the CNN the "bare" Green's function without any information about the sign. Surprisingly, as illustrated in the inset of Fig. 4, the state function F for the phase-multiplied Green's functions does not exhibit a notable improvement in predicting the position of the phase transition over the bare Green's function. While the function moves slightly in parameter space, it also acquires a much broader spread
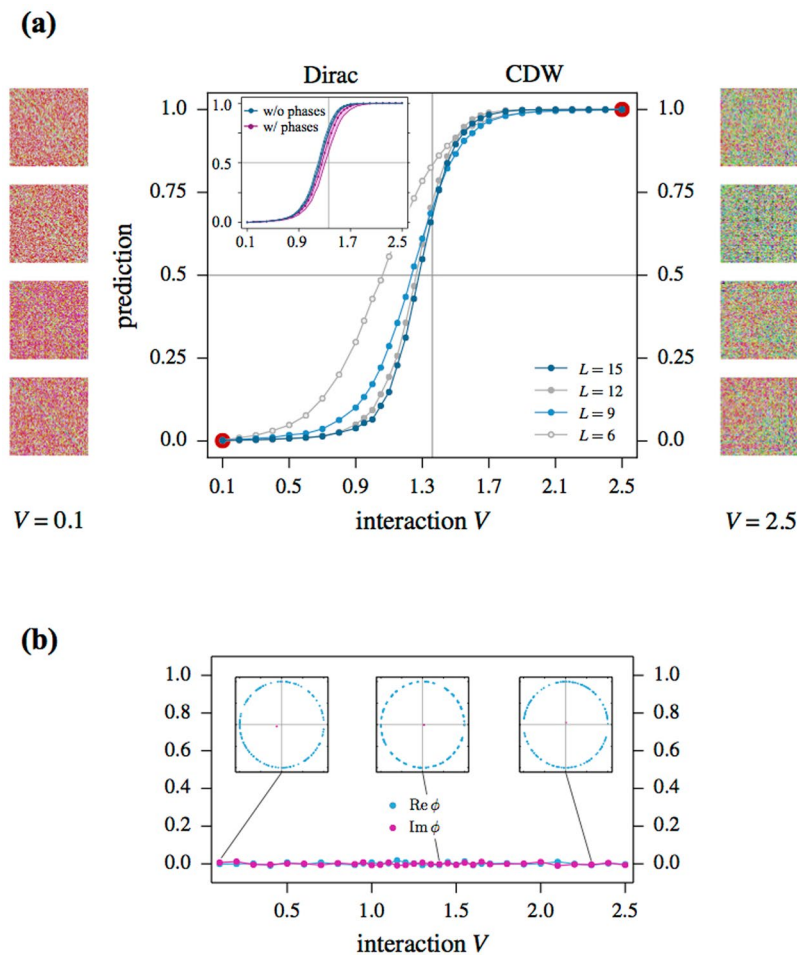
Figure 4. (**a**) Prediction of a CNN for the phase transition from a Dirac semi-metal to a charge density wave (CDW) ordered state in the half-filled spinless fermion Hubbard model (10) on the honeycomb lattice of size $2 \cdot L \times L$. The CNN has been trained on 8192 representative samples of the bare Green's function deep inside the two phases (indicated by the red dots). The images in the left and right columns are color-converted instances of the Green's function used in the training. The inset shows a comparison of the prediction for the $L = 9$ system when feeding the CNN with the bare Green's function or the Green's function multiplied by the relative sign/complex phase associated with each configuration (of a given Markov chain). (**b**) The averaged real and imaginary part of the weight's phase ϕ, $\mathrm{Re}(\phi)$ and $\mathrm{Im}(\phi)$, respectively, is shown in the main part of the figure for $L = 6$. The three insets show the distribution of the phase for a sequence of 128 measurements, with their average depicted by the pink dot.

(estimated from averaging over 12 epochs, see the Methods section)[49]. Considering the data for different system sizes in Fig. 4 one can determine a quantitative estimate of the location of the fermionic phase transition, which is in very good agreement with the Monte Carlo results[5, 44]. This convincingly demonstrates that the CNN is capable of providing a high-quality state function F discriminating the two fermionic phases, even when the sign content of the configurations is ignored. Importantly, we note that the approach with bare Green's matrices can provide a significant gain in computational efficiency over that which includes information about the relative sign of individual configurations, by sampling multiple parallel Markov chains. Thus, in light of the results of Fig. 4 (inset), which show no systematic improvement of the state function F given additional information on the sign structure, we choose to show results for the bare Green's functions in the examples below. The fact that such an approach produces a highly accurate state function F is a striking demonstration of the power of QMC + machine learning, even in models afflicted with a serious sign problem.

Next, we consider the spinless fermion system of Eq. (10) at one-third filling. Going below half-filling turns the itinerant phase for small coupling V into a conventional metal with a nodal Fermi line, while for large V we still expect some sort of CDW-ordered Mott insulating state. In contrast to half-filling, the one-third-filled system has no known sign-free (Majorana) basis. Applying our QMC + machine learning approach to this problem, we again find that a state discriminating function F can be identified by a properly optimized CNN. This procedure indicates the existence of a phase transition around $V_c \approx 0.7 \pm 0.1$ as illustrated in Fig. 5, which matches a recent estimate from entanglement calculations[48]. The precise nature of the Mott insulating phase at large V has so far
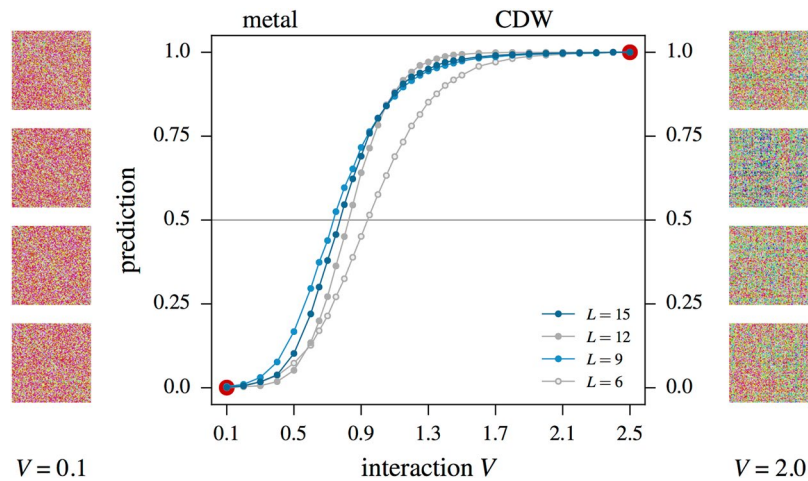
Figure 5. CNN-based identification of the phase transition in the one-third filled, spinless Hubbard model (10) on the honeycomb lattice with nearest-neighbor repulsion V. The side panels show representative samples of the Green's functions at the two reference points $V = 0.1$ and $V = 2.5$. The network finds two clearly separated phases of finite extent. For $V \lesssim 0.7$ a metallic phase is realized, which at weak coupling is known to exhibit a nodal line Fermi surface, while for $V \gtrsim 0.7$ the system forms a charge density wave (CDW).
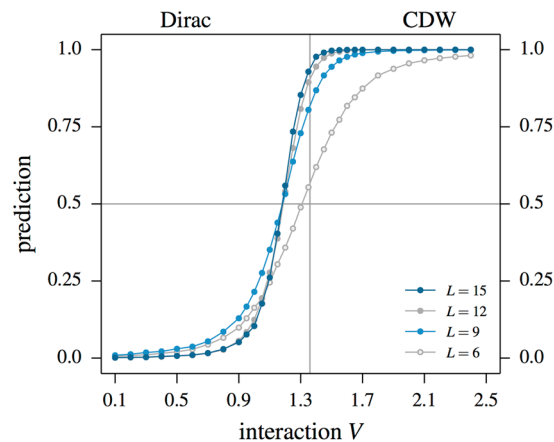
Figure 6. An example of transfer learning in an artificial neural network. A CNN that was trained to discriminate the phases of the sign-problem free, *spinful* Hubbard model (4) and then applied to identify the phases and phase transition of the sign-problematic, *spinless* Hubbard model (10). The network is found to reliably distinguish the fermionic phases of the spinless model and provides a relatively accurate estimate for the location of the phase transition (the vertical line indicates the location of the transition in the thermodynamic limit of infinite system size).

remained elusive, which unfortunately is not altered by the supervised learning approach employed in the current study.

Finally, we explore whether we can generate "transfer learning" by training a neural network on one model, then using the trained network to discriminate phases from configurations produced for an entirely different Hamiltonian. This approach was highly successful for neural networks trained with classical Ising configurations in ref. 14. Here, using samples of the Green's function, we train a CNN to discriminate the fermionic phases of the sign-problem free, spinful fermion model (4) and then apply the trained network for supervised learning on the sign-problematic, spinless fermion model (10). This procedure seems justified based on the fact that at half-filling the two models exhibit similar physics, with the potential energy driving a Gross-Neveu type phase transition from a Dirac semi-metal to a SDW/CDW charge-ordered phase, respectively. Results for the predictions of the averaged state function are illustrated in Fig. 6, which shows that the CNN is capable of reliably distinguishing the fermionic phases of the spinless model, even producing a rough estimate for the location of the phase transition. Thus, we find that this approach indeed allows for a certain level of transfer learning between sign-problem free and sign-problematic Hamiltonians, suggesting a fruitful area of future study on the relationship between supervised machine learning and the sign problem.

Discussion

We have introduced a powerful numerical scheme to reliably distinguish fermionic phases of matter by a combination of quantum Monte Carlo sampling and a subsequent machine learning analysis of the sampled Green's functions via a convolutional neural network. Our numerical experiments for a family of Hubbard-type models demonstrate that this approach extends to sign-problematic many-fermion models that are not amenable to the conventional QMC approach of sampling and statistical analysis. These findings thereby provide a perspective on the information content of the sampled *ensemble* of Green's functions. In contrast to a conventional statistical physics analysis, in which equal-time correlation functions calculated from this ensemble of Green's functions exhibit a statistical uncertainty so large that they are rendered completely unusable, the machine learning approach demonstrates that the same ensemble of Green's functions holds sufficient information to positively discriminate fermionic phases. This Green's function based machine learning approach is very general and can be applied to QMC flavors beyond the auxiliary field techniques applied in the current work. In particular, this approach can be readily adapted by world-line Monte Carlo approaches which are highly successful in the study of quantum magnets and bosonic systems. For the future, we envision to refine and improve our "phase recognition" machine learning approach such that it can complement the existing statistical analysis of QMC data in mapping out phase diagrams of quantum many-body systems.

Methods

Machine learning.   Neural networks come in a huge variety of different architectures; precisely which setup to choose for a specific problem is answered by selecting the empirically most successful architecture. In this paper, we started with a setup, see Fig. 1, that is successfully used to classify images such as the CIFAR-10 dataset[50]. Its network architecture consists of two main components – a convolutional and a fully connected part. The convolutional part processes the data by a combination of two convolutional and max pooling units. Both of these units are activated by a rectified linear function (relu) and have filters of size 3×3. The total number of filters is 32 for the first and 64 for the second. The data is then fed into a fully connected, relu activated layer of 512 neurons. To avoid overfitting, we applied a dropout regularization at a rate of 0.5 to this layer. At the output of the CNN we consider a fully connected softmax layer. The optimization of the neural network is performed using a cross entropy as a cost function and ADAM[51] as a particularly efficient variant of the stochastic gradient at a learning rate of $\gamma = 0.0001$. The network was trained over 16 epochs and results were averaged over the last 8 epochs. Our numerical implementation of the neural network is based on the TensorFlow library[52].

Location of the phase transition.   An alternative choice to identify the location of the phase transition for a finite system size is the parametric location of the inflection point of the prediction function *F*. At the inflection point the derivative is maximal, i.e. moving the coupling parameters slightly towards any one of the two phases leads to the largest possible change in the prediction. Such an indicator of the phase transition is commonly used e.g. when considering susceptibility measurements of conventional order parameters. For our setup we find almost no distinguishable difference in the location and finite-size scaling of the inflection point and the point at which the prediction function $F = 1/2$.

Determinantal Quantum Monte Carlo.   For our DQMC simulation, we use a projective algorithm with a discretization step of $\Delta\tau = 0.1$ and a projection time $\theta = 10$. Thus, the auxiliary field for the spinful Hubbard model is of size $2 \cdot L^2 \times 200$. The Green's functions are of size $2 \cdot L^2 \times 2 \cdot L^2$. The test wave function $|\psi_T\rangle$ is generated by taking the quadratic part of the Hamiltonian and randomizing the hopping strengths strongly enough such that the eigenvalues of adjacent states are separated by more than 10^{-3}. The eigenvectors corresponding to the lowest $N_{\text{particles}}$ eigenvalues are used for the test wave function.

The Hubbard-Stratonovich transformation is applied to each quartic operator, introducing the auxiliary field. For the models studied in this paper, one such transformation has to be carried out for each site or each bond, respectively, and on each slice in projection time. There are various ways to perform this transformation, in particular, one is often free to choose the channel one performs the transformation in and what type of field should be created. One possible realization is to decouple a density-density interaction of strength *U* with general indices α and β denoting for example spin and/or lattice site in the following way

$$e^{-\Delta\tau U n_\alpha n_\beta} = \frac{1}{2} \sum_{s=\pm 1} \prod_{a=\alpha,\beta} e^{-\left(s\lambda + \frac{U\Delta\tau}{2}\right)\left(n_a - \frac{1}{2}\right)}.$$

(11)

where the auxiliary variable *s* is in $\{\pm 1\}$ and λ is a constant related to *U*. This transformation results in complex weights for $U > 0$, i.e. a repulsive interaction. In the spinful Hubbard model at half filling, the product of the phases of the two determinants and the prefactor result in an overall prefactor of 1 for the weight, i.e. there is no sign problem. This changes drastically once one moves away from half filling or takes away one of the fermion species, resulting in a severe sign or phase problem. An alternative transformation that allows us to work with real numbers only works by decoupling in the magnetization channel. While at first look computationally favorable (because of the real numbers), it turns out that the convergence of magnetic observables is significantly better in the complex case, as it retains the $SU(2)$ symmetry explicitly while in the real case this symmetry is only restored after the summation over all configurations has been carried out.

For the phase sensitive calculations, one can in principle calculate the absolute phase of a weight from the determinant in Eq. (9). However, this approach is found to be plagued by numerical instabilities making its computation prohibitively expensive in terms of computing resources. Alternatively, one may track the changes in the phase along the Markov chain and thus calculate the *relative* phase with respect to an initial phase for each configuration visited in the Markov chain. The change in phase ϕ'/ϕ is given by the phase of the ratio of weights

$W(C')/W_C$ between the current configuration C and a proposed configuration C'. Using this quantity, the initial phase ϕ is updated by multiplying ϕ with the ratio of phases for adjacent steps on the Markov chain

$$\phi \xrightarrow{\cdot \frac{\phi'}{\phi}} \phi' \xrightarrow{\cdot \frac{\phi''}{\phi'}} \dots \xrightarrow{\cdot \frac{\phi^{n+1}}{\phi^n}} \phi^{n+1}. \tag{12}$$

Using the relative phase has the advantage that it is possible to compute this quantity with very high accuracy, while it is not expected to change any of the physics (a global transformation of the phase of the weights is compensated when normalizing the partition or wave function).

References

1. Hirsch, J. E., Sugar, R. L., Scalapino, D. J. & Blankenbecler, R. Monte carlo simulations of one-dimensional fermion systems. *Phys. Rev. B* **26**, 5033 (1982).
2. Loh, E. Y. *et al.* Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B* **41**, 9301 (1990).
3. Chandrasekharan, S. & Wiese, U.-J. Meron-Cluster Solution of Fermion Sign Problems. *Phys. Rev. Lett.* **83**, 3116 (1999).
4. Chandrasekharan, S. & Li, A. Fermion bag solutions to some sign problems in four-fermion field theories. *Phys. Rev. D* **85**, 091502 (2012).
5. Li, Z.-X., Jiang, Y.-F & Yao, H. Solving the fermion sign problem in quantum Monte Carlo simulations by Majorana representation. *Phys. Rev. B* **91**, 241117 (2015a).
6. Bravyi, S., DiVincenzo, D., Oliveira, R. & Terhal, B. The complexity of stoquastic local Hamiltonian problems. *Quant. Inf. Comp.* **8**, 0361 (2008).
7. Kaul, R. K., Melko, R. G. & Sandvik, A. W. Bridging Lattice-Scale Physics and Continuum Field Theory with Quantum Monte Carlo Simulations. *Annual Review of Condensed Matter Physics* **4**, 179 (2013).
8. Troyer, M. & Wiese, U.-J. Computational Complexity and Fundamental Limitations to Fermionic Quantum Monte Carlo Simulations. *Phys. Rev. Lett.* **94**, 170201 (2005).
9. Zaanen, J. *et al.* Towards a complete theory of high Tc. *Nature Phys.* **2**, 138 (2006).
10. Schofield, A. J. Non-Fermi liquids. *Contemp. Phys.* **40**, 95 (1999).
11. Löhneysen, H. V., Rosch, A., Vojta, M. & Wölfle, P. Fermi-liquid instabilities at magnetic quantum phase transitions. *Mod. Phys.* **79**, 1015 (2007).
12. Balents, L. Spin liquids in frustrated magnets. *Nature* **464**, 199 (2010).
13. Landau, D. P. & Binder, K. *A Guide to Monte Carlo Simulations in Statistical Physics* (Cambridge University Press, 2000).
14. Carrasquilla, J. & Melko, R. G. Machine learning phases of matter. *Nature Physics* **13**, 431 (2017).
15. Wang, L. Discovering phase transitions with unsupervised learning. *Phys. Rev. B* **94**, 195105 (2016).
16. Schmidhuber, J. Deep learning in neural networks: An overview. *Networks* **61**, 85 (2015).
17. Nielsen, M. A. *Neural Networks and Deep Learning* (Determination Press, 2015).
18. Deng, D.-L., Li, X. & Sarma, S. D. Exact machine learning topological states, arXiv:1609.09060.
19. Tubman, N. M. Measuring quantum entanglement, machine learning and wave function tomography: Bridging theory and experiment with the quantum gas microscope, arXiv:1609.08142.
20. Tanaka, A. & Tomiya, A. Detection of phase transition via convolutional neural network. *J. Phys. Soc. Jpn.* **86**, 063001 (2017).
21. van Nieuwenburg, E. P., Liu, Y.-H. & Huber, S. D. Learning phase transitions by confusion. *Nature Physics* **13**, 435 (2017).
22. Liu, J., Qi, Y., Meng, Z. Y. & Fu, L. Self-learning monte carlo method. *Phys. Rev. B* **95**, 041101 (2017).
23. Huang, L. & Wang, L. Accelerate monte carlo simulations with restricted boltzmann machines. *Phys. Rev. B* **95**, 035105 (2017).
24. Zhang, Y. & Kim, E.-A. Triangular quantum loop topography for machine learning. *Phys. Rev. Lett.* **118**, 216401 (2017).
25. Portman, N. & Tamblyn, I. Sampling algorithms for validation of supervised learning models for Ising-like systems. arXiv:1611.05891.
26. Huang, L., Feng Yang, Y. & Wang, L. Recommender engine for continuous time quantum monte carlo methods. *Phys. Rev. E* **95**, 031301 (2017).
27. Ohtsuki, T. & Ohtsuki, T. Deep learning the quantum phase transitions in random electron systems: Applications to three dimensions. *J. Phys. Soc. Jpn.* **86**, 044708 (2017).
28. Biamonte, J. *et al.* Quantum machine learning, arXiv:1611.09347.
29. Bény, C. Deep learning and the renormalization group, arXiv:1301.3124.
30. Mehta, P. & Schwab, D. J. An exact mapping between the Variational Renormalization Group and Deep Learning, arXiv:1410.3831.
31. Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
32. Schoenholz, S. S., Cubuk, E. D., Sussman, D. M., Kaxiras, E. & Liu, A. J. A structural approach to relaxation in glassy liquids. *Nat. Phys.* **12**, 469 (2016).
33. Amin, M. H., Andriyash, E., Rolfe, J., Kulchytskyy, B. & Melko, R. Quantum Boltzmann Machine, arXiv:1601.02036.
34. Torlai, G. & Melko, R. G. Learning thermodynamics with boltzmann machines. *Phys. Rev. B* **94**, 165134 (2016).
35. Carleo, G. & Troyer, M. Solving the Quantum Many-Body Problem with Artificial Neural Networks. *Science* **355**, 602 (2017).
36. Stoudenmire, E. & Schwab, D. J. In Advances in Neural Information Processing Systems 29, edited by Lee, D. D., Sugiyama, M., Luxburg, U. V., Guyon, I. & Garnett, R., pp. 4799–4807 (Curran Associates, Inc., 2016).
37. Otsuka, Y., Yunoki, S. & Sorella, S. Universal quantum criticality in the metal-insulator transition of two-dimensional interacting dirac electrons. *Phys. Rev. X* **6**, 011029 (2016).
38. Hirsch, J. E. Discrete hubbard-stratonovich transformation for fermion lattice models. *Phys. Rev. B* **28**, 4059 (1983).
39. Duchon, E., Loh, Y. L. & Trivedi, N. In Novel Superfluids: Volume 2, edited by Bennemann, K.-H. and Ketterson, J. B., Chap. 17 (Oxford University Press, Oxford, 2014).
40. Ch'ng, K., Carrasquilla, J., Melko, R. G. & Khatami, E. Machine Learning Phases of Strongly Correlated Fermions, arXiv:1609.02552.
41. Foley, J. D., van Dam, A., Feiner, S. K. & Hughes, J. F. *Computer Graphics: Principles and Practice (2nd Ed.)* (Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1990).
42. Meng, Z. Y., Lang, T. C., Wessel, S., Assaad, F. F. & Muramatsu, A. Quantum spin liquid emerging in two-dimensional correlated Dirac fermions. *Nature (London)* **464**, 847 (2010).
43. Gross, D. J. & Neveu, A. Dynamical symmetry breaking in asymptotically free field theories. *Phys. Rev. D* **10**, 3235 (1974).
44. Wang, L., Corboz, P. & Troyer, M. Fermionic quantum critical point of spinless fermions on a honeycomb lattice. *New Journal of Physics* **16**, 103008 (2014).
45. Li, Z.-X., Jiang, Y.-F. & Yao, H. Fermion-sign-free Majarana-quantum-Monte-Carlo studies of quantum critical phenomena of Dirac fermions in two dimensions. *New Journal of Physics* **17**, 085003 (2015b).
46. Motruk, J., Grushin, A. G., de Juan, F. & Pollmann, F. Interaction-driven phases in the half-filled honeycomb lattice: An infinite density matrix renormalization group study. *Phys. Rev. B* **92**, 085147 (2015).

47. Capponi, S. & Läuchli, A. M. Phase diagram of interacting spinless fermions on the honeycomb lattice: A comprehensive exact diagonalization study. *Phys. Rev. B* **92**, 085146 (2015).
48. Broecker, P. & Trebst, S. Entanglement and the fermion sign problem in auxiliary field quantum Monte Carlo simulations. *Phys. Rev. B* **94**, 075144 (2016).
49. Multiplying the Green's matrices by a completely *random* phase is found to give a similar shift and broadening of the spread.
50. Krizhevsky, A. Learning Multiple Layers of Features from Tiny Images. *Tech. Rep*. (2009).
51. Kingma, D. & Ba, J. Adam: A Method for Stochastic Optimization, arXiv:1412.6980.
52. Abadi, M. *et al*. TensorFlow: Large-Scale Machine Learning on Heterogeneous Distributed Systems, arXiv:1603.04467; Software available from tensorflow.org.

Acknowledgements

Author Contributions

P.B. developed the ideas underlying this manuscript in discussion with all authors. P.B. and J.C. implemented all simulations codes, P.B. collected the numerical data. Data analysis and interpretation was jointly done by all authors. All authors contributed to the writing of the manuscript.

Additional Information

Supplementary information accompanies this paper at doi:10.1038/s41598-017-09098-0

Competing Interests: The authors declare that they have no competing interests.

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