Editors' Suggestion

# Learnability scaling of quantum states: Restricted Boltzmann machines

Dan Sehayek,[1,2] Anna Golubeva,[1,2] Michael S. Albergo,[1,2] Bohdan Kulchytskyy,[1,2] Giacomo Torlai,[3] and Roger G. Melko[1,2]

[1]*Department of Physics and Astronomy, University of Waterloo, Ontario, Canada N2L 3G1*

[2]*Perimeter Institute for Theoretical Physics, Waterloo, Ontario, Canada N2L 2Y5*

[3]*Center for Computational Quantum Physics, Flatiron Institute, New York, New York 10010, USA*

(Received 26 August 2019; revised manuscript received 23 October 2019; published 15 November 2019)

Generative modeling with machine learning has provided a new perspective on the data-driven task of reconstructing quantum states from a set of qubit measurements. As increasingly large experimental quantum devices are built in laboratories, the question of how these machine learning techniques scale with the number of qubits is becoming crucial. We empirically study the scaling of restricted Boltzmann machines (RBMs) applied to reconstruct ground-state wave functions of the one-dimensional transverse-field Ising model from projective measurement data. We define a learning criterion via a threshold on the relative error in the energy estimator of the machine. With this criterion, we observe that the number of RBM weight parameters required for accurate representation of the ground state in the worst case – near criticality – scales quadratically with the number of qubits. By pruning small parameters of the trained model, we find that the number of weights can be significantly reduced while still retaining an accurate reconstruction. This provides evidence that overparametrization of the RBM is required to facilitate the learning process.

DOI: 10.1103/PhysRevB.100.195125

## I. INTRODUCTION

Generative models are a powerful class of machine learning algorithms that seek to reconstruct an unknown probability distribution $p(\mathbf{x})$ from a set of data $\mathbf{x}$. After training, generative models can be used to estimate the likelihood of new data not contained in the original set or to produce new data samples for various purposes. Recently, industry-standard generative models were repurposed by the physics community with the goal of reconstructing a quantum wave function from projective measurement data [1–3]. The question of scalability is of paramount importance for the reconstruction of quantum states prepared by near-term hardware which comprises tens or hundreds of qubits.

While several generative modeling techniques are available for quantum state reconstruction, by far the most well studied involves restricted Boltzmann machines (RBMs) [1,2,4–6]. RBMs can be used to explicitly parametrize a probability distribution $p(\mathbf{x})$ and, through a suitable complex generalization, a quantum wave function [2,7]. One main application of RBMs is the data-driven reconstruction of experimental states, which was recently demonstrated for a Rydberg-atom quantum simulator [8]. These and other uses have been covered extensively in the literature, including several recent reviews [9–11].

With the steady increase in the size of experimental quantum devices, an important question is how data-driven quantum state reconstruction scales with the number of qubits. While many results have been reported for fixed finite-size reconstructions, less work has been done in the way of scaling analyses [12]. Particularly important is the difference in scaling complexity of approximate machine learning methods for practical reconstructions, compared to full quantum state tomography that, in general, scales exponentially [13].

In this paper, we present a systematic study of the scaling of the computational resources required for accurate reconstruction of a quantum state. In particular, we focus on RBMs used to reconstruct the ground-state wave function of a one-dimensional transverse-field Ising model, which has a positive-real representation. Our training data are a set of projective measurements sampled independently from a simulated tensor-network wave function. We define a learning criterion based on the accuracy of the energy estimator of the RBM. The state reconstruction is considered successful when the relative error of the energy estimator is smaller than a fixed threshold. We target in particular two contributions to the asymptotic scaling behavior in the many-qubit limit: the representational power of the neural network, i.e., the *expressiveness* of the parametrization of the state, and the amount of data required to train the model, also known as the *sample complexity*.

We find that deep within the ferromagnetic and paramagnetic phases, the number of RBM parameters required for accurate representation of the ground state is $O(1)$. As the transverse field is varied to approach the quantum critical point between these two phases, the state becomes more challenging to reconstruct, as expected due to long-range quantum correlations that arise there. At the critical point, we observe that under standard RBM training procedures the number of parameters grows quadratically in the number of qubits, $O(N^2)$. The minimum number of measurements required to train this number of parameters scales linearly with the number of qubits, $O(N)$. Interestingly, we find that the number of parameters required for an accurate reconstruction can be significantly reduced posttraining by pruning small weights and fine-tuning the RBM by a small number of additional training iterations. We argue that an RBM requires

 ©2019 American Physical Society

overparameterization to facilitate the optimization procedure associated with learning.

## II. DEFINING A SCALING STUDY

We are interested in probing the asymptotic scaling of the computational resources required to reconstruct a quantum state using an RBM. The training set comprises projective measurement data produced from the ground-state wave function of the one-dimensional transverse-field Ising model (TFIM) defined by the Hamiltonian

$$H = -J \sum_{\langle ij \rangle} \sigma_i^z \sigma_j^z - h \sum_i \sigma_i^x, \tag{1}$$

where $\sigma^{x,y,z}$ are Pauli operators, defined over $N$ sites (or qubits), and $\langle ij \rangle$ denotes nearest-neighbor pairs on a one-dimensional lattice with open boundary conditions. This model is thoroughly studied in the condensed-matter and quantum information literature and serves as a standard benchmark for many numerical methods, such as quantum Monte Carlo [14,15], tensor networks (TNs) [16], and more recent quantum optimization algorithms [17–19]. We generate training data from a density matrix renormalization group (DMRG) simulation [20] for various values of $h/J$ using the ITENSOR library [21]. The measurements of the ground-state wave function are produced in the $\sigma^z$ basis.

The Perron-Frobenius theorem guarantees that when the Hamiltonian (1) has negative off-diagonal matrix elements in the $\sigma^z$ (computational) basis, the ground-state wave function is positive real. Thus, there is a direct mapping between the wave function and a probability distribution, $\psi(\boldsymbol{\sigma}) = \sqrt{p(\boldsymbol{\sigma})}$. This allows for a significant simplification in the RBM network structure since complex phases or signs need not be parametrized. In addition, the computational basis is trivially informationally complete, enabling training from data produced only in the $\sigma^z$ basis [2].

### A. Restricted Boltzmann machine

The RBM consists of two layers of binary variables $v_i, h_j \in \{0, 1\}$. The energy associated with each configuration is given by

$$E_{\boldsymbol{\lambda}}(\boldsymbol{v}, \boldsymbol{h}) = -\sum_{ij} W_{ij} v_i h_j - \sum_i^N b_i v_i - \sum_j^{N_h} c_j h_j, \tag{2}$$

where $N$ is the number of visible units, representing the qubits or spins, and $N_h$ is the number of hidden units parametrizing the interactions. The two layers are fully connected via the weight matrix $\boldsymbol{W}$ that, along with the bias terms $b_i$ and $c_j$, forms the set of learnable parameters $\boldsymbol{\lambda} = (\boldsymbol{W}, \boldsymbol{b}, \boldsymbol{c})$. The energy function (2) defines the joint probability distribution

$$p_{\boldsymbol{\lambda}}(\boldsymbol{v}, \boldsymbol{h}) = \frac{1}{Z_{\boldsymbol{\lambda}}} e^{-E_{\boldsymbol{\lambda}}(\boldsymbol{v}, \boldsymbol{h})}, \tag{3}$$

where $Z_{\boldsymbol{\lambda}}$ is the partition function of the machine. The marginal distribution is obtained by tracing out the hidden units,

$$p_{\boldsymbol{\lambda}}(\boldsymbol{v}) = \sum_{\boldsymbol{h}} p_{\boldsymbol{\lambda}}(\boldsymbol{v}, \boldsymbol{h}) = \frac{1}{Z_{\boldsymbol{\lambda}}} \sum_{\boldsymbol{h}} e^{-E_{\boldsymbol{\lambda}}(\boldsymbol{v}, \boldsymbol{h})}. \tag{4}$$

It is this marginal distribution that forms the approximate representation of the ground state, $\psi_{\boldsymbol{\lambda}}(\boldsymbol{v}) = \sqrt{p_{\boldsymbol{\lambda}}(\boldsymbol{v})}$. In other words, because of the assumed positive-real form of the wave function, the training procedure is equivalent to conventional unsupervised learning of an RBM [22]. In particular, the objective of the training procedure is to minimize the Kullback-Leibler (KL) divergence, which defines the discrepancy between the distribution of projective measurements and the probability distribution parameterized by the RBM, through a method known as contrastive divergence [23]. In the present work, we use the QUCUMBER software package to implement and train a positive-real RBM [24].

### B. Learning criterion

In order to quantify the resources required for the data-driven reconstruction of the ground-state wave function for the TFIM, one must be able to assess when the learning is "complete." Generally, the fidelity is considered a standard measure of the closeness of two quantum states, such as a target state and an approximate reconstructed state. However, in more generic situations than ours, where a TN representation of the target quantum state may not be available, calculations of the fidelity typically scale exponentially, which renders them intractable for even moderate numbers of qubits. An alternative method for defining the accuracy of a reconstruction is to measure expectation values of local observables. Such expectation values can be efficiently calculated through standard estimators from samples produced by the RBM. Importantly, these can be compared with the exact values measured from our DMRG simulations or through other methods such as quantum Monte Carlo that do not admit an explicit representation of the ground state.

The relative error between an RBM estimator and an exact DMRG expectation value will be referred to as the relative observable error (ROE). For the current study, we define the *learning criterion* through the ROE in the expectation value of the energy, which can be calculated from the RBM using standard Markov chain Monte Carlo techniques. Take $\bar{U} = \langle H \rangle_{\text{RBM}}$ to be the average of the energy estimator calculated from $n$ samples generated by the RBM. Since $n$ is finite, a statistical error exists in the estimator, quantified by the standard deviation $\sigma$. To account for this in a relative error measure, we compute the Gaussian confidence interval given by $\bar{U} \pm C \frac{\sigma}{\sqrt{n}}$. The value of $C = 2.576$ corresponding to 99% confidence will be used throughout this paper. If $U = \langle H \rangle_{\text{exact}}$ is the exact value of the energy estimator (calculated, e.g., with DMRG), then we can upper bound the ROE by the larger relative error value of the confidence interval:

$$\epsilon = \max \left| \frac{U - (\bar{U} \pm C\sigma\sqrt{n})}{U} \right|. \tag{5}$$

Essentially, this means that we consider the learning to be complete when our desired upper bound on the ROE is satisfied 99% of the time on our sample size. We find empirically that $\epsilon = 0.002$ is a reasonable value that can be achieved by RBMs trained on TFIM data with conventional algorithms for $N \leqslant 100$ qubits. At smaller values (e.g., $\epsilon = 0.001$) training becomes impractical for $N > 50$, while for larger values we observed that the results reported below remain qualitatively

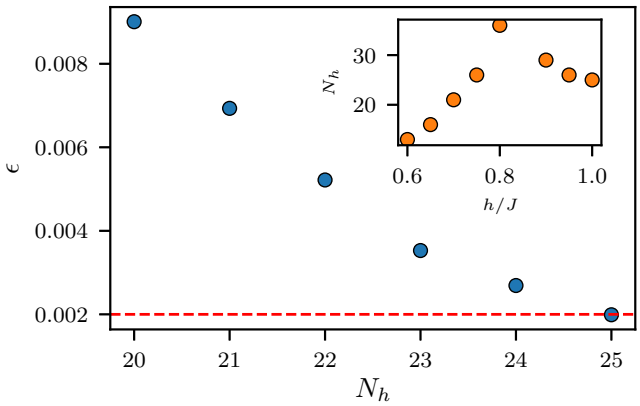


FIG. 1. The procedure used to determine the RBM expressiveness required to represent the TFIM wave function at $h/J = 1$ with $N = 50$ qubits. The number of hidden units $N_h$ is increased until the desired $\epsilon$ is achieved. The inset illustrates the number of hidden units required for convergence to $\epsilon \leqslant 0.002$ for different values of $h/J$ near criticality. The position of the peak is discussed in the text.

the same; thus, we use $\epsilon = 0.002$ in the remainder of the paper.

With this learning criterion, we analyze the scaling behavior of the RBM by controlling two variables: the number of model parameters per qubit and the number of training measurements $M$, i.e., the sample complexity. However, we note that, as is typical in machine learning studies, many other variables exist that are related to network architecture, learning rates, batch size, etc., referred to as *hyperparameters*. Here these hyperparameters have been made consistent for all values of $h/J$ and all system sizes $N$.

## III. RESULTS

In this section we present numerical results for the scaling of computational resources for reconstruction of the TFIM ground-state wave function for several values of $h/J$. In order to systematically investigate scaling, we control variables of interest in different ways, as described in the following sections.

### A. Scaling of the model parameters

To begin, we are interested in the minimal number of RBM parameters per qubit required to faithfully reproduce the ground-state energy. We parametrize this with the scaling of the size of the hidden layer $N_h$. We consider the critical point, corresponding to $h/J = 1$, as well as the ferromagnetic and paramagnetic phases. For each value of $N$, we produce large numbers of projective measurements of $\sigma^z$ values using the DMRG simulation of the TFIM. Then, effectively assuming that the number of available training samples $M \to \infty$, we increase the number of hidden units $N_h$ until the learning criterion is uniquely satisfied for each value of $N$.

Our procedure is illustrated in Fig. 1 for a fixed system size of $N = 50$. In the main plot, corresponding to $h/J = 1$, we observe that the specified learning criterion $\epsilon = 0.002$ cannot be achieved for $N_h < 25$. The minimum number of parameters required to accurately represent the ground-state

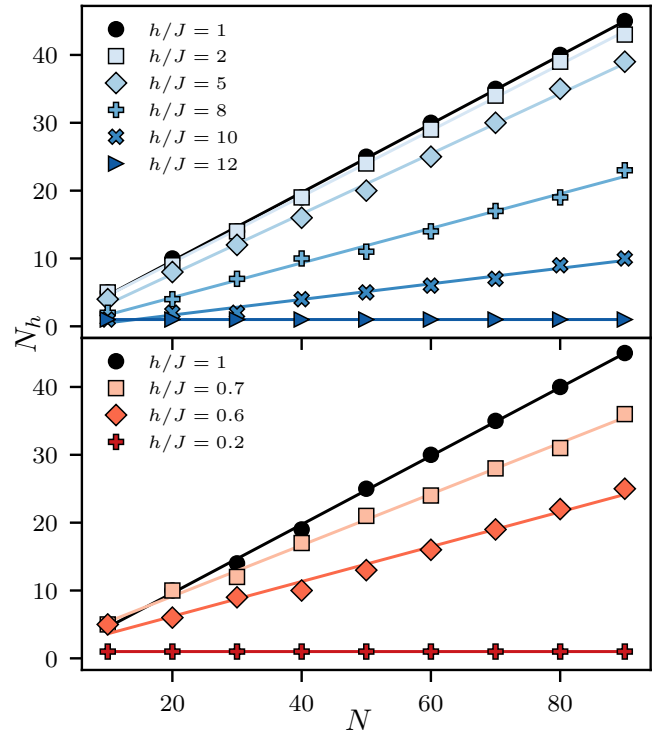


FIG. 2. Minimum number of hidden units $N_h$ required for $\epsilon \leqslant 0.002$ for various values of $h/J$. Straight lines are fits to the data.

wave function is thus $N_h = 25$. The inset illustrates the dependence of $N_h$ on field values near $h/J = 1$, where the quantum wave function is most entangled. One would expect that in the limit $N \to \infty$ the number of parameters required to accurately parametrize the wave function would be maximal at $h/J = 1$. Curiously, we find that this peak occurs around $h/J \approx 0.8$, slightly on the ferromagnetic side from the critical point. We hypothesize that this feature might be tied to the magnetization of the underlying data set used for training, which was produced by our DMRG simulations in ITENSOR. For the maximum bond dimension that we employ (2000), the expected $\mathbb{Z}_2$ symmetry is not realized below certain values of the transverse field strength $h$ when the number of qubits is large. Furthermore, a similar phenomenon has been observed previously in studies of the relative energy in diffusion Monte Carlo [15] and a recent variational imaginary time ansatz [19]. It would be an interesting topic of future study.

The result of repeating the above procedure for various numbers of qubits $N$ is illustrated in Fig. 2. For values of $h/J$ deep within the ferromagnetic or the paramagnetic phase, the required minimum number of hidden units scales as $N_h \sim O(1)$ in the asymptotic limit of large $N$. This reflects the informational simplicity of the data set close to the ferromagnetic or paramagnetic limits. Near $h/J \approx 1$, the scaling is clearly linear, $N_h \sim N$, meaning that the leading-order scaling of the number of parameters is $O(N^2)$, as each additional hidden unit quadratically scales the number of elements in the weight matrix $\boldsymbol{W}$. Due to the presence of the bias terms in Eq. (2), we would also expect a subleading term that scales proportionally to $N$; however, as noted in the Appendix, for data sets with an underlying $\mathbb{Z}_2$ symmetry, these bias terms do not represent independent parameters for the purpose of wave

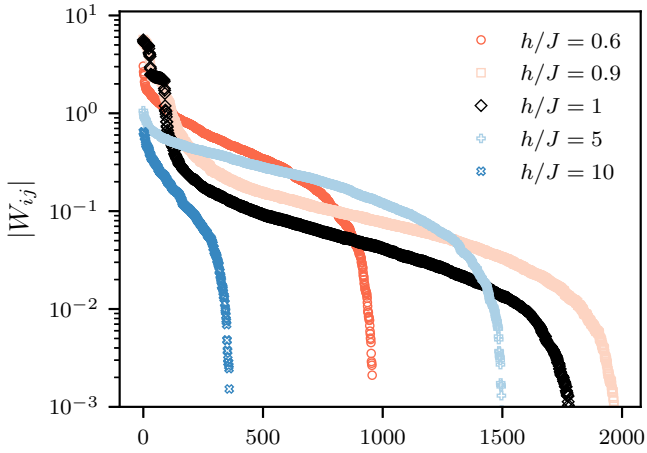


FIG. 3. Weight magnitudes, sorted in descending order from left to right, for various transverse field values and $N = 60$. Converged RBM models from the parameter study shown in Fig. 2 are used here.

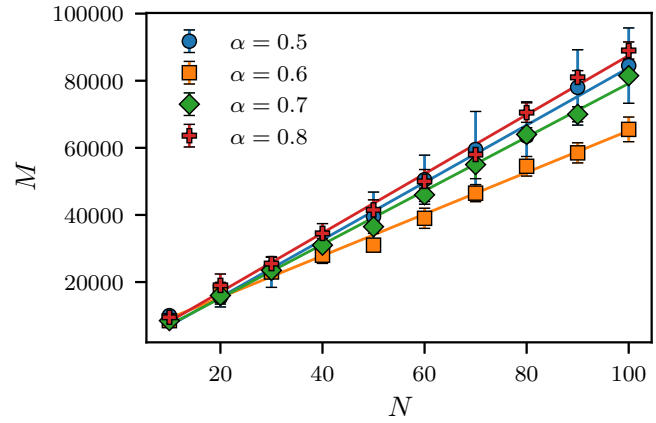


FIG. 4. The minimum number of training examples $M$ required for $\epsilon \leqslant 0.002$ for the TFIM at the critical point $h/J = 1$ for different ratios of the number of hidden to visible units $\alpha = N_h/N$.

function reconstruction. Finally, we note that for larger ROE thresholds $\epsilon > 0.002$ the prefactors and slopes are different, but the asymptotic scaling of the number of hidden units still remains linear near criticality.

We draw a brief comparison between the above results and naive expectations for the scaling of the number of parameters in the simplest TN representation of a one-dimensional wave function: the matrix product state (MPS). We refer the reader to recent reviews on the topic, e.g., Refs. [25,26]. For a finite-size MPS with $N$ matrices, the simplest estimate for the number of parameters required to store a wave function is $O(N\chi^2)$, where $\chi$ is the *bond dimension*. The bond dimension required for good representation accuracy depends on the amount of entanglement in the system. In the presence of a bipartition, the entanglement entropy $S$ is upper bounded by the logarithm of the rank of the reduced density matrix. For an MPS, every bond that lies on the bipartition therefore gives an entropy contribution of, at most, $\ln(\chi)$, leading to the scaling rule of $\chi \sim \exp(S)$. For the one-dimensional TFIM, we expect $S \sim O(1)$ in the ordered phases and $S \sim \ln(N)$ at the quantum critical point. This yields a naive scaling of $O(N)$ in the ordered case and $O(N^3)$ at the critical point. We note that our RBM scaling is more consistent with a translationally invariant MPS encoding, where these asymptotic scaling complexities are both reduced by a factor of $N$. This reduction is expected in long one-dimensional chains, where tensors sufficiently far from the edges are typically identical.

Further insight on the RBM result can be obtained by examining the distribution of the absolute weight values $|W_{ij}|$ in a typical trained model. In Fig. 3 we plot the magnitude of each individual weight, sorted in decreasing order from left to right, on a logarithmic scale. One can see that near criticality, the largest contribution is given by the first 10%–20% of weights; then the weight values decrease exponentially in magnitude, eventually falling off even more rapidly. We return to this observation after the next section.

### B. Scaling of sample complexity

Above, we studied the minimum number of RBM parameters required to find an accurate ground-state energy, assuming access to an infinite amount of training data. We now determine the minimal sample complexity required to accurately train this number of parameters. We focus on $h/J = 1$ and fix the ratio $\alpha = N_h/N$ for several values near $1/2$. Then, repeating the procedure from the last section, we increment the number of training examples $M$ by 2500 until the ROE learning criterion $\epsilon \leqslant 0.002$ on the RBM energy estimator is achieved. This procedure is repeated for a number of different initial weight configurations, and the results are averaged. The resulting scaling of the sample complexity is shown in Fig. 4.

The results suggest that for $N_h/N$ near $1/2$, the sample complexity scales linearly in the number of qubits. Combining the asymptotic scaling results from the previous two sections, $N_h \sim N$ and $M \sim N$, suggests that the number of samples per parameter required to train a minimally expressive machine scales as $N/N^2 \sim 1/N$. In other words, the relative "data cost" required to train a new weight parameter decreases with an increasing number of qubits. A linear scaling of the sample complexity was also observed in a recent generative modeling scheme based on positive operator-valued measurements [3].

We remark that a sample complexity linear in $N$ is consistent with observations on the probably approximately correct (PAC) learnability of quantum states. In Ref. [27] Aaronson argues that, if one is concerned only about learning a state well enough to predict the outcomes of most measurements drawn from it, the exponential cost usually associated with full state tomography is reduced to a linear scaling in $N$. This is what we find in Fig. 4. Indeed, a characteristic of the Aaronson learning theorem is the assumption that the training samples are drawn *independently* from the probability distribution. This is exactly the setting that we employ in training the RBM in the present work. Hence, it is reasonable to expect the theorem to apply.

### C. Reducing the number of model parameters posttraining

We now return to the results of Sec. III A, where it was found that the minimal number of hidden units required to satisfy our chosen ROE scales approximately as $N_h \approx \frac{1}{2}N$ near the TFIM quantum critical point. Implicit in this result is the RBM optimization procedure used to train the machine: a

TABLE I. The number of weights required to achieve $\epsilon \leqslant 0.002$ at the critical point $h/J = 1$. Results for the "original" RBM are taken from Fig. 2.

| | $N$ | | | |
|---|---|---|---|---|
| | 10 | 20 | 30 | 40 |
| Original | 50 | 200 | 420 | 760 |
| Pruned | 20 | 50 | 79 | 119 |

stochastic gradient descent that minimizes the KL divergence [9]. Since it is not obvious that the scaling behavior is independent of this optimization procedure, it is fair to ask the question: Is it possible to find more efficient representations—with fewer model parameters—by modifying the learning protocol? Indeed, it is known that the required number of model parameters is intertwined with the specifics of the training procedure. In particular, it has been found that the overparameterization inherent in deep neural networks can ease and accelerate their optimization by (stochastic) gradient descent [28–32].

Figure 3 offers a clue that the RBM parametrization may not be optimal (i.e., minimal) for the final trained wave function by demonstrating that the distribution of the weight magnitudes in a trained model is very nonuniform: 10% to 20% of the weights have values that are orders of magnitude larger than the rest. Recent machine learning literature has studied the relative importance of these smaller weights with a procedure called *pruning*. Following the ideas of Refs. [33,34], we define a pruning procedure for our scaling study in the following steps:

(1) Start from the original, converged trained model (e.g., Fig. 3, with $N_h = \frac{1}{2}N$ for $h/J = 1$).

(2) Set a threshold $\delta$ for the weight magnitudes. If a given $|W_{ij}| < \delta$, set $W_{ij} = 0$, and freeze it for the following steps.

(3) Fine-tune the pruned model by running several more training iterations until the desired accuracy (as defined by the ROE learning criterion) is restored.

(4) Repeat steps 2–4, pruning additional weights until the model fails to fulfill the learning criterion.

We choose the pruning threshold such that 40% of the nonzero weights are pruned in the first iteration, and 5% of the nonzero weights are pruned in each following iteration. Note also that in this procedure we do not prune biases (see the Appendix for further comments).

We apply weight pruning to our trained RBM focusing on the critical point of the TFIM and find that a significant reduction in the number of RBM parameters required to correctly capture the critical TFIM ground-state energy can be achieved for all system sizes. The results for several small numbers of qubits are presented in Table I. We interpret this to mean that the standard training of an RBM with contrastive divergence benefits from an overparameterization, employing more weights than is strictly required for accurate expression of the TFIM wave function in order to make the optimization more navigable. We note that in some rare cases, pruning a very small number of weights seriously alters the ROE, highlighting that some paths through the optimization landscape may depend on weight parameters that are not redundant. For this reason, rigorous uncertainty intervals on our results are difficult to estimate at present.

The success of the pruning procedure opens up the possibility of systematically searching for a change in scaling behavior. However, due to the significant increase in methodological complexity introduced by the pruning procedure, this analysis is out of scope for the current study and will be presented in another work.

## IV. DISCUSSION

In this paper, we have empirically studied the scaling of computational resources required for the accurate reconstruction of positive-real wave functions using generative modeling with a restricted Boltzmann machine. We obtained scaling results by examining the energy estimator calculated from an RBM after training on projective measurement data from the one-dimensional transverse-field Ising model (TFIM) ground state. An RBM reconstruction of the ground-state wave function was defined to be "accurate" when the relative error between the RBM estimator and the exact energy value was below a fixed threshold. Thus, scaling results in this paper are subject to the caveat that they could change if other criteria were to be considered, such as the convergence of fidelity or correlation functions.

In the present case, convergence of the relative error in the energy produces several interesting results. First, for a standard optimization procedure with contrastive divergence, the number of weights required for accurate reconstruction is, at best, constant (deep in the ferromagnetic/paramagnetic phases). At worst, this scaling is quadratic; this occurs near the quantum critical point between the two phases. It is interesting to note that such scaling is consistent with that expected from a translationally invariant matrix product state encoding. In addition, the minimum number of samples required to converge the energy at the critical point is observed to scale linearly with the number of qubits. This is consistent with a theorem by Aaronson that predicts a linear scaling in a similar setting for PAC learning [27].

Further, we present evidence that the number of parameters required to represent the ground state is drastically affected by the RBM learning procedure. By employing a pruning technique that sets small weights to zero, then fine-tuning the remaining model parameters through additional training, we observe a very significant reduction in the number of parameters required to accurately reproduce the energy. It would be interesting to examine whether the asymptotic scaling functions identified here are affected by the pruning procedure. Further, such a technique could provide a systematic way of searching for the minimal model expressiveness required for a given quantum state. It would then be interesting to compare the obtained results to theoretical expectations for the representational capacity of RBMs required for quantum ground-state wave functions [4,35,36].

Indeed, numerous recent results have highlighted the benefit that overparameterization provides for optimizing deep learning models [28,31,32]. In this paper, we have discovered that RBMs trained on measurement data for positive-real wave functions may as well be aided by overparameterization beyond what is needed for the theoretical representation of

the quantum state as a means of assisting the standard optimization procedure of minimizing the KL divergence via contrastive divergence. The question of how to systematically mitigate this overparameterization while still maintaining the ease of optimization is an active area of research [37–39], one whose successes will be of great use for more efficiently representing and studying quantum systems.

It is natural to wonder what the scaling of computational resources is for reconstructing quantum states that are not real or positive. To this end, the present results point towards a rich field of similar scaling studies that should be pursued on a variety of quantum many-body models in the future. The question of scaling is also especially pertinent for state-of-the-art experiments, such as fermionic quantum simulators [40], wave functions generated by quantum dynamics [41,42], and quantum chemistry calculations with superconducting circuits [43]. In contrast to positive wave functions, the reconstruction (with a suitably modified RBM) demands training data from an extended set of measurement bases. The ability to theoretically identify the minimal set, and how the size of this set scales with the number of qubits, will ultimately determine the feasibility of integrating this type of machine learning technology into such near-term quantum devices.

In conclusion, we have proposed a systematic procedure to evaluate the scaling of resources for reconstructing positive-real wave functions with RBMs. A tighter threshold in the reconstructed energy accuracy, or improved neural-network parametrizations of nonpositive states, will likely require a more powerful breed of generative model. Recurrent neural networks, transformers, and other autoregressive models are currently being considered in this context. In light of the fact that current intermediate-scale quantum devices are already capable of producing training data on tens and even hundreds of qubits, we expect these and similar scaling studies will be pursued in earnest in the near future.

## ACKNOWLEDGMENTS

We acknowledge enlightening discussions with M. Beach, J. Carrasquilla, I. De Vlugt, M. Ganahl, E. Inack, D. Kong, E. Merali, A. Rochetto, and D. Sels. The DMRG calculations were performed using the ITENSOR libraries [21]. The RBM calculations were performed with the QUCUMBER package [24]. This work was made possible by the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET) and Compute Canada. A.G. is supported by NSERC. R.G.M. is supported by NSERC, the Canada Research Chair program, and the Perimeter Institute for Theoretical Physics. Research at the Perimeter Institute is supported in part by the government of Canada through the Department of Innovation, Science and Economic Development Canada and by the province of Ontario through the Ministry of Economic Development, Job Creation and Trade. The Flatiron Institute is supported by the Simons Foundation.

## APPENDIX

Projective measurement data in the $\sigma^z$ basis for the TFIM will generally obey $\mathbb{Z}_2$ symmetry in the absence of symmetry breaking, which may occur due to a limited DMRG bond dimension (or can happen spontaneously in the thermodynamic limit). Let us assign the measurement on a single qubit $\sigma_i^z = \sigma_i = \pm 1$. The probability of any given state over $N$ qubits $\boldsymbol{\sigma}$ is therefore the same as that of the corresponding spin-flipped state $\bar{\boldsymbol{\sigma}}$; that is, the magnetization of the state will be zero. In a typical RBM, the values of the visible and hidden units will be 0 and 1, which can be mapped to an "occupation" number rather than a spin magnetization. One can always consider instead the spin language by working in the $\pm 1$ basis. In this basis, when the underlying data set has zero magnetization, we can assume that the energy of the RBM takes the form

$$E_{\boldsymbol{\lambda}}(\boldsymbol{\sigma}^v, \boldsymbol{\sigma}^h) = -\sum_{ij} \tilde{W}_{ij} \sigma_i^v \sigma_j^h, \tag{A1}$$

where $\sigma_i^v$ ($\sigma_j^h$) is a single visible (hidden) unit in the spin language. It can then be shown that this energy function, when used to define a joint distribution, results in $p_{\boldsymbol{\lambda}}(\boldsymbol{\sigma}^v) = p_{\boldsymbol{\lambda}}(\bar{\boldsymbol{\sigma}}^v)$ after marginalizing over the hidden units. In other words, the RBM in the $\pm 1$ representation requires no biases (or magnetic fields) to capture a $\mathbb{Z}_2$ invariance in the data distribution.

To see what this means in the occupation number representation, one can map $\sigma_i^v = 2v_i - 1$ and $\sigma_j^h = 2h_j - 1$, where $v_i$ and $h_j \in \{0, 1\}$. By transforming the energy expression (A1) and setting $W_{ij} = 4\tilde{W}_{ij}$, $b_i = -\sum_j 2\tilde{W}_{ij}$, and $c_j = -\sum_i 2\tilde{W}_{ij}$, we obtain an expression identical to Eq. (2). This allows us to interpret the presence of biases, which are learned by the RBM even in TFIM data sets that are observed to be $\mathbb{Z}_2$ invariant.

Let us examine the weight matrix of a typical converged run of the RBM. From the above arguments, we can calculate the ratios of the biases to the sums of weights along rows and columns of the weight matrix:

$$\alpha_i = \sum_j W_{ij}/b_i = -\sum_j 4\tilde{W}_{ij} \Big/ \sum_j 2\tilde{W}_{ij} = -2, \tag{A2}$$

$$\beta_j = \sum_i W_{ij}/c_j = -\sum_i 4\tilde{W}_{ij} \Big/ \sum_i 2\tilde{W}_{ij} = -2. \tag{A3}$$

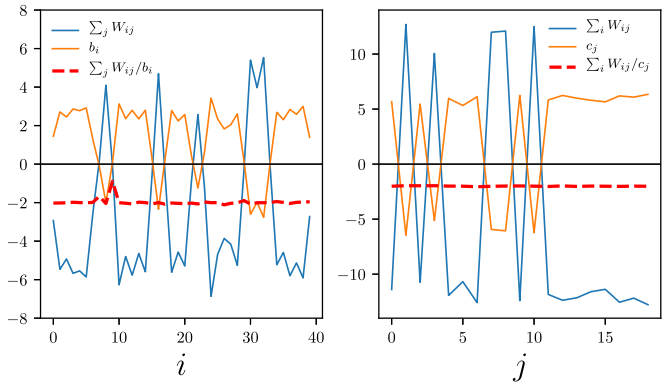


FIG. 5. The left and right panels show the values of the weight matrix from an RBM trained on $N = 40$ at the quantum critical point to $\epsilon \leqslant 0.002$, summed along its rows and columns, respectively. One can see that the ratios of the summed weights to biases fluctuate around $-2$, as predicted.

Thus, one expects at least approximately that $\alpha_i = \beta_j = -2$ for all $i$ and $j$ for any of the trained RBM models considered in Fig. 2. We confirm this behavior for $N = 40$ and $N_h = 20$ in Fig. 5.


[1] G. Torlai and R. G. Melko, Phys. Rev. B **94**, 165134 (2016).
[2] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo, Nat. Phys. **14**, 447 (2018).
[3] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Nat. Mach. Intell. **1**, 155 (2019).
[4] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Phys. Rev. B **97**, 085104 (2018).
[5] S. Cheng, J. Chen, and L. Wang, Entropy **20**, 583 (2018).
[6] G. Carleo, Y. Nomura, and M. Imada, Nat. Commun. **9**, 5322 (2018).
[7] G. Carleo and M. Troyer, Science **355**, 602 (2017).
[8] G. Torlai, B. Timar, E. P. L. van Nieuwenburg, H. Levine, A. Omran, A. Keesling, H. Bernien, M. Greiner, V. Vuletić, M. D. Lukin, R. G. Melko, and M. Endres, arXiv:1904.08441.
[9] G. Torlai and R. G. Melko, arXiv:1905.04312.
[10] R. G. Melko, G. Carleo, J. Carrasquilla, and J. I. Cirac, Nat. Phys. **15**, 887 (2019).
[11] G. Carleo, I. Cirac, K. Cranmer, L. Daudet, M. Schuld, N. Tishby, L. Vogt-Maranto, and L. Zdeborová, arXiv:1903.10563.
[12] S. Czischek, J. M. Pawlowski, T. Gasenzer, and M. Garttner, arXiv:1907.12844.
[13] *Quantum State Estimation*, edited by M. Paris and J. Rehacek, Lecture Notes in Physics Vol. 649 (Springer, Berlin, 2004).
[14] A. W. Sandvik, Phys. Rev. E **68**, 056701 (2003).
[15] E. M. Inack, G. Giudici, T. Parolini, G. Santoro, and S. Pilati, Phys. Rev. A **97**, 032307 (2018).
[16] G. Vidal, Phys. Rev. Lett. **99**, 220405 (2007).
[17] W. W. Ho and T. H. Hsieh, SciPost Phys. **6**, 29 (2019).
[18] W. W. Ho, C. Jonay, and T. H. Hsieh, Phys. Rev. A **99**, 052332 (2019).
[19] M. J. S. Beach, R. G. Melko, T. Grover, and T. H. Hsieh, Phys. Rev. B **100**, 094434 (2019).
[20] A. J. Ferris and G. Vidal, Phys. Rev. B **85**, 165146 (2012).
[21] ITENSOR library, version 2.0.11, http://itensor.org.
[22] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, Cognit. Sci. **9**, 147 (1985).
[23] G. E. Hinton, Neural Comput. **14**, 1771 (2002).
[24] M. J. S. Beach, I. D. Vlugt, A. Golubeva, P. Huembeli, B. Kulchytskyy, X. Luo, R. G. Melko, E. Merali, and G. Torlai, SciPost Phys. **7**, 9 (2019).
[25] R. Orús, Ann. Phys. (NY) **349**, 117 (2014).
[26] R. Orús, Nat. Rev. Phys. **1**, 538 (2019).
[27] S. Aaronson, Proc. R. Soc. A **463** (2007).
[28] D. Lopez-Paz and L. Sagun, Easing non-convex optimization with neural networks, *International Conference on Learning Representations (ICLR)* (2018).
[29] R. Livni, S. Shalev-Shwartz, and O. Shamir, Advances in Neural Information Processing systems conference (NIPS), arXiv:1410.1141 (2014).
[30] S. Arora, N. Cohen, and E. Hazan, International Conference on Machine Learning (ICML), arXiv:1802.06509 (2018).
[31] Z. Allen-Zhu, Y. Li, and Z. Song, International Conference on Machine Learning (ICML), arXiv:1811.03962 (2019).
[32] K. A. Sankararaman, S. De, Z. Xu, W. R. Huang, and T. Goldstein, arXiv:1904.06963.
[33] S. Han, J. Pool, J. Tran, and W. Dally, in *Advances in Neural Information Processing systems conference (NIPS)*, edited by C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett (Curran Associates, 2015), pp. 1135–1143.
[34] D. C. Mocanu, E. Mocanu, P. H. Nguyen, M. Gibescu, and A. Liotta, Mach. Learn. **104**, 243 (2016).
[35] X. Gao and L.-M. Duan, Nat. Commun. **8**, 662 (2017).
[36] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Phys. Rev. X **8**, 011006 (2018).
[37] J. Frankle and M. Carbin, International Conference on Learning Representations (ICLR), arXiv:1803.03635 (2019).
[38] H. Zhou, J. Lan, R. Liu, and J. Yosinski, Advances in Neural Information Processing systems conference (NeurIPS), arXiv:1905.01067 (2019).
[39] N. Lee, T. Ajanthan, and P. H. S. Torr, International Conference on Learning Representations (ICLR), arXiv:1810.02340 (2019).
[40] A. Mazurenko, C. S. Chiu, G. Ji, M. F. Parsons, M. Kanász-Nagy, R. Schmidt, F. Grusdt, E. Demler, D. Greif, and M. Greiner, Nature (London) **545**, 462 (2017).
[41] B. P. Lanyon, C. Maier, M. Holzäpfel, T. Baumgratz, C. Hempel, P. Jurcevic, I. Dhand, A. S. Buyskikh, A. J. Daley, M. Cramer, M. B. Plenio, R. Blatt, and C. F. Roos, Nat. Phys. **13**, 1158 (2017).
[42] A. Keesling, A. Omran, H. Levine, H. Bernien, H. Pichler, S. Choi, R. Samajdar, S. Schwartz, P. Silvi, S. Sachdev, P. Zoller, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, Nature **568**, 207 (2019).
[43] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, Nature (London) **549**, 242 (2017).