

Machine learning phase transitions with a quantum processorA. V. Uvarov^{ⓧ,*}, A. S. Kardashin,[†] and J. D. Biamonte^{ⓧ‡}*Deep Quantum Laboratory, Skolkovo Institute of Science and Technology, 3 Nobel Street, Moscow 143026, Russia*

(Received 24 June 2019; revised 31 March 2020; accepted 19 June 2020; published 15 July 2020)

Machine learning has emerged as a promising approach to unveil properties of many-body systems. Recently proposed as a tool to classify phases of matter, the approach relies on classical simulation methods—such as Monte Carlo—which are known to experience an exponential slowdown when simulating certain quantum systems. To overcome this slowdown while still leveraging machine learning, we propose a variational quantum algorithm which merges quantum simulation and quantum machine learning to classify phases of matter. Our classifier is directly fed labeled states recovered by the variational quantum eigensolver algorithm, thereby avoiding the data-reading slowdown experienced in many applications of quantum enhanced machine learning. We propose families of variational ansatz states that are inspired directly by tensor networks. This allows us to use tools from tensor network theory to explain properties of the phase diagrams the presented quantum algorithm recovers. Finally, we propose a majority vote quantum classifier built from a nearest-neighbor (checkerboard) quantum neural network. This classifier is successfully trained to recognize phases of matter with 99% accuracy for the transverse field Ising model and 94% accuracy for the XXZ model. These findings suggest that a merger between quantum simulation and quantum enhanced machine learning offers a fertile ground to develop computational insights into quantum systems.

DOI: [10.1103/PhysRevA.102.012415](https://doi.org/10.1103/PhysRevA.102.012415)**I. INTRODUCTION**

The best contemporary algorithms to emulate quantum systems using classical computers suffer from an exponential slowdown in limiting cases. A recent approach is to apply machine learning, which offers new techniques for large-scale data analysis. In particular, machine learning was recently proposed as a tool to recognize phases of matter [1,2]. These methods still rely on Monte Carlo sampling (or alternative classical simulation methods) which suffers from the an exponential slowdown induced by the so-called sign problem. Independently, quantum algorithms have also been proposed as a platform for machine learning [3–11]. In addition, unlike classical algorithms, quantum simulators are predicted to simulate quantum systems efficiently [12–14]. Here we merge quantum machine learning with quantum simulation, leveraging quantum mechanics to overcome two classical bottlenecks. Namely, we leverage quantum algorithms as a tool to simulate quantum systems by preparing states which are labeled and fed into a quantum classifier. The latter removes the data-reading slowdown experienced in many applications of quantum enhanced machine learning, while the former utilizes quantum simulators to avoid classical methods with known limitations. We also replace the standard unitary coupled cluster ansatz found in implementations of the variational quantum eigensolver [15] with families of tensor network ansatz states.

Using machine learning techniques, one can analyze the phase diagrams of strongly interacting quantum systems and thus directly address system properties. Monte Carlo samples of such systems are fed as input data and classified by using supervised [1,2] or unsupervised [16] learning. Hamiltonians with up to a few hundreds of spins can be studied. Nonetheless, for fermionic systems, the use of Monte Carlo methods is drastically restricted by the sign problem.

In the variational quantum circuits approach [17–19], a quantum processor is required to prepare a sufficiently rich variety of probe states. The variational quantum eigensolver (VQE) represents an implementation of variational quantum circuits which uses a quantum processor to prepare a family of states characterized by a polynomial number of parameters to minimize the expectation value of a given Hamiltonian [15,20,21].

In this paper, we circumvent the sign problem by using a quantum variational algorithm realizable on near-term quantum computers. To classify the phases of a given quantum Hamiltonian, we first prepare its approximate ground states variationally, and then feed them as an input to a quantum classifier. In this respect, there is no need to sample microscopic configurations with Monte Carlo based methods. Instead, the classifier has direct access to the quantum states [7], yielding thus an effective realization of quantum-enhanced machine learning.

II. TENSOR NETWORK ANSATZ STATES FOR VQE

VQE is a hybrid iterative quantum-classical algorithm used to approximate the ground state of a given Hamiltonian [15]. It relies on preparing an ansatz state $|\psi(\theta)\rangle$ by applying a sequence of quantum gates $U(\theta)$ and sampling the expectation

*alexey.uvarov@skoltech.ru

†andrey.kardashin@skoltech.ru

‡j.biamonte@skoltech.ru

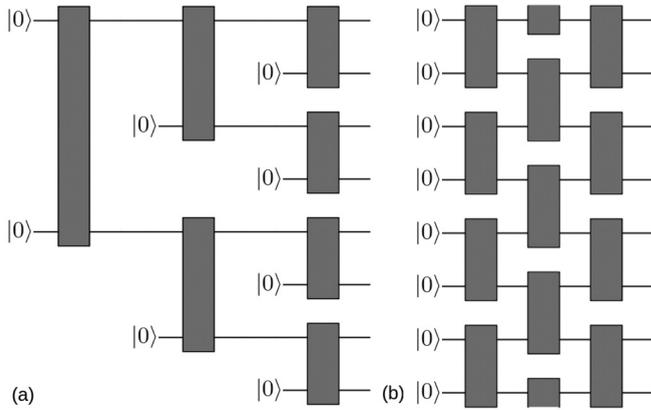


FIG. 1. (a) Tree tensor network state. (b) Checkerboard tensor network state. In both cases, the quantum register is instantiated in the $|0\rangle^{\otimes n}$ state and subject to entangling gates. Black boxes indicate two-qubit gates specified by Eq. (2) and Fig. 3. Each block is parametrized independently. Herein quantum circuits are read left to right.

value of a given Hamiltonian relative to this state, followed by a classical optimizer to minimize the energy, $\langle \psi(\theta) | H | \psi(\theta) \rangle$. Within the VQE method, we approach the ground state of a given Hamiltonian by using tensor network ansatz states.

We proceed by representing the Hamiltonian as a sum of tensor products of Pauli operators:

$$H = \sum \mathcal{J}_{\alpha_1 \alpha_2 \dots \alpha_n} \sigma_{\alpha_1} \otimes \sigma_{\alpha_2} \otimes \dots \otimes \sigma_{\alpha_n}, \quad (1)$$

where $\alpha_i \in \{0, 1, 2, 3\}$ enumerate the Pauli matrices $\{\mathbb{1}, X, Y, Z\}$. With the decomposition shown in Eq. (1), individual terms of $\langle \psi(\theta) | H | \psi(\theta) \rangle$ can be estimated and variationally minimized elementwise by using a classical-to-quantum process. In each iteration one prepares the state $|\psi(\theta)\rangle$ and measures each qubit in the local X , Y , or Z basis, estimates the energy, and updates θ . This method can become scalable if the number of terms in the Hamiltonian is polynomially bounded in the number of spins and the coefficients $\mathcal{J}_{\alpha_1 \alpha_2 \dots \alpha_n}$ are defined up to poly(n) digits.

The performance of VQE crucially depends on the choice of the ansatz state. A common approach is to use the unitary version of the coupled cluster method, the UCC ansatz [20,22–24]. For interacting spin problems, the (non)unitary coupled cluster ansatz can be composed out of spin-flip operators [17,25]. There is no known classical algorithm to efficiently implement this method, even when the series is truncated to low-order terms [26]. In principle, a quantum computer could efficiently prepare this state, truncated up to some k th order by using the Suzuki–Trotter decomposition [27]. However, for a system of n qubits it requires $O(n^k)$ unitary gates, making this technique out of reach for the available quantum computers. Still, even if UCC is truncated to single and double interactions (UCCSD), it requires $9n^2$ operations and necessitates applying some optimization strategy [28].

Instead, we test a number of shorter ansatz states inspired by tensor network states, namely, (i) a rank-one circuit; (ii) a tree tensor network circuit Fig. 1(a); and (iii) a family of checkerboard-shaped circuits [Fig. 1(b)] with varying depth.

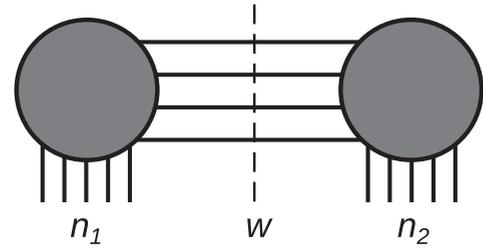


FIG. 2. A quantum circuit can be treated as a tensor network state with all bonds having dimension two. If a bipartition cuts w wires, the total bond dimension is at most 2^w , while the cut separates at most w ebits of entanglement.

These states differ in the amount of entanglement they can support. In general, quantum states of n qubits can be represented by n -index tensors, while quantum circuits are embodied by tensor networks. Each quantum gate is seen as a vertex, and each string is an index running through $\{0, 1\}$, while the maximum amount of entanglement is determined by the number of strings one needs to cut to separate the subsystems (see Fig. 2). Each string corresponds to at most one ebit of entanglement. An n -qubit state can contain at most $\lfloor n/2 \rfloor$ ebits of entanglement. To formalize the latter, suppose there exists a certain bipartition in the system that brings n_1 qubits to the first subsystem and n_2 qubits to the second. It is then possible to regroup the tensor network state into a $2^{n_1} \times 2^{n_2}$ matrix. The rank of this matrix provides an upper bound to the amount of entanglement across this bipartition: a rank- k state can support at most $\ln_2 k$ ebits of entanglement, i.e., when it is in the maximally entangled state.

III. STATE PREPARATION

We first approximate the ground state of a Hamiltonian by rank-one states. One can prepare any unentangled state using $2n$ gates by subsequently applying R_y and R_z rotations to each qubit. This ansatz captures the first-order truncation of UCC.

Figure 1(a) illustrates a circuit implementing a *tree tensor network state*. This state is prepared by using the subsequently described two-qubit parametric blocks. The preparation procedure is easier to explain when the number of qubits is a power of two. A block first entangles qubits 1 and $(n/2 + 1)$. In the next layer, two blocks act on qubit pairs $(1, n/4 + 1)$ and $(n/2 + 1, n/2 + n/4 + 1)$. Then, for each half of the register, we act in the same way, but now the blocks act on qubit pairs $(1, n/4 + 1)$ and $(n/2 + 1, n/2 + n/4 + 1)$. Each half of the register is again divided into two halves. The pattern continues. For systems where the number of qubits is not a power of two, one can do this procedure up until the number of qubits involved is the closest power of two and then make the last layer of operators incomplete. Such a structure enables long-range correlations but limits the entanglement entropy that a state bipartition can potentially have: any contiguous region can be isolated by $O(1)$ cuts. It is easy to contract $\langle \psi_{\text{tree}} | A | \psi_{\text{tree}} \rangle$ classically with A being a local observable and $|\psi_{\text{tree}}\rangle$ being a tree tensor network state.

There is some freedom in choosing the two-qubit blocks that comprise the tree tensor network. In principle, one can implement any unitary in $SU(4)$ by using three

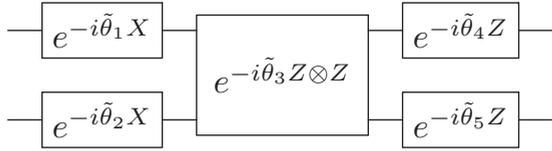


FIG. 3. Two-qubit entangler gate used in preparation of the states.

controlled-NOTs and 15 single-qubit rotations [29]. However, throughout this work we used two-qubit gates with fewer parameters $\tilde{\theta}$ to simplify the optimization (Fig. 3):

$$U(\tilde{\theta}) = [R_z(\tilde{\theta}_5) \otimes R_z(\tilde{\theta}_4)] \circ R_{zz}(\tilde{\theta}_3) \circ [R_x(\tilde{\theta}_1) \otimes R_x(\tilde{\theta}_2)], \quad (2)$$

where $R_z(\theta) = e^{-i\theta Z}$, $R_x(\theta) = e^{-i\theta X}$, $R_{zz}(\theta) = e^{-i\theta Z \otimes Z}$. Thus, a complete ansatz would have five free parameters per two-qubit block. In total, the tree tensor network ansatz features $n - 1$ blocks, yielding $5n - 5$ independent parameters.

Remarkably, the block shown in Fig. 2 is inspired by the parametrized Hamiltonian approach [30] and the unitary operators used in the quantum approximate optimization algorithm (QAOA) [31]. Of course, the ansatz with such blocks is less expressive than if each block implemented the entire SU(4) group. However, such an ansatz could also have some redundancy as the ansatz gates are applied to a fixed n -qubit input state $|0\rangle^{\otimes n}$.

In a *checkerboard tensor network*, the entangling blocks are positioned in a checkerboard pattern as shown in Fig. 1(b). In the following, we impose periodic boundary conditions, meaning that the last qubit is linked to the first. For this ansatz, we also use the two-qubit entangling gate shown in Fig. 3. The ansatz has $5L\lfloor n/2 \rfloor$ independent parameters, where L is the number of layers in the circuit.

To isolate a region in a checkerboard tensor network state with d layers, one has to cut at least $O(d)$ bonds regardless of the region size [if the region is very small, one can also make a “horizontal cut” of $O(L)$ bonds]. Therefore, if we set the number of layers to be equal to $\lfloor \ln_2(n) \rfloor$, the upper bound on the entanglement scaling is equal to that in critical one-dimensional systems. To implement any maximally entangled state; that is, a state with the maximum possible amount of ebits, one needs to cut $\lfloor \frac{n}{2} \rfloor$ bonds. If the checkerboard ansatz has open boundary conditions, it needs at least $2\lfloor \frac{n}{2} \rfloor$ gates to saturate entanglement. Periodic boundary conditions make the ansatz more powerful and lower this bound by half to $\lfloor \frac{n}{2} \rfloor$ layers.

IV. QUANTUM CLASSIFIER

Not only can the checkerboard tensor network be used as a VQE ansatz but it also functions as a quantum majority vote classifier. Each data point is a VQE solution: the parameters of the unitary gates are optimized to obtain the minimum energy $\langle \psi(\theta) | H(h) | \psi(\theta) \rangle$, where h determines the phase of the model. Each VQE solution is labeled with “0” or “1” depending on whether the model parameter is above or below the phase-transition point. We then prepare a circuit made of two parts (Fig. 4). The first part takes the blank qubit registry $|000\dots 0\rangle$ and prepares the VQE solution in the form of an ansatz state. The second part takes this state as an input and applies a unitary $U_{\text{class}}(\varphi)$. We then measure the output of the

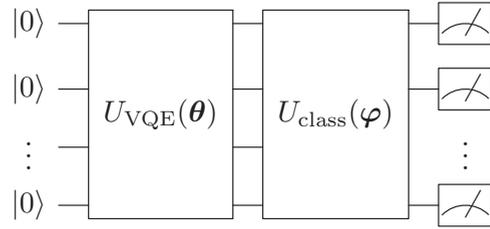


FIG. 4. Quantum circuit that implements the classifier. The first part prepares the VQE solution, the second one performs the classification. The assigned label is inferred from the measurements in the Z basis. Both U_{VQE} and U_{class} have the checkerboard structure.

circuit in the Z basis. Let q_0 and q_1 be the total number of measurements in which more than half of the qubits are in “0” or “1” states, respectively. The classifier returns the predicted probability $p = q_1/(q_0 + q_1)$ for the state belonging to the class “1” being equal to the probability that the majority of qubits vote “1,” excluding ties. The quantum circuit is shown in more detail in Appendix D.

Let $\{(\theta_i, y_i)\}_{i=1}^{N_{\text{train}}}$ be the set of training data points and their labels $y_i \in \{0, 1\}$. Let $p_i \in [0, 1]$ be the label predicted by the neural network. Then the loss function is

$$f = - \sum_{i=1}^{N_{\text{train}}} [y_i \ln p_i + (1 - y_i) \ln(1 - p_i)]. \quad (3)$$

minimize f , we used the simultaneous perturbation stochastic approximation (SPSA) algorithm [32]. This algorithm estimates the gradient vector by computing a finite difference in a random direction, then performs a gradient descent step. We optimized the log loss over 300 epochs, with both finite difference step size and learning rate starting very coarse and decreasing as $1/\sqrt{n_{\text{epoch}}}$, where n_{epoch} is the epoch number.

V. NUMERICAL RESULTS

We apply our quantum circuit of $n = 10$ qubits to the transverse-field Ising model (TFIM), while being exactly solvable [33,34] this model serves for testing purposes. This model is specified by the Hamiltonian

$$H_{\text{TFIM}} = J \sum_{i=1}^n \sigma_i^z \sigma_{i+1}^z + h \sum_{i=1}^n \sigma_i^x, \quad J > 0, \quad h > 0, \quad (4)$$

where σ_i^α is the Pauli matrix α acting on the i th spin, and we impose periodic boundary conditions $\sigma_{n+1}^\alpha = \sigma_1^\alpha$.

The ground state of TFIM is determined by the trade-off between Heisenberg exchange coupling, the first term in Eq. (4), favoring collinear orientation of magnetic moments in the z direction, and the Zeeman coupling to the transverse magnetic field, the second term. The latter has the tendency to flip the z component, being thus the source of “quantum fluctuations” in the system. In a magnetically ordered state, a strong magnetic field $h \geq J$ destroys magnetic order even at zero temperature. This induces quantum fluctuations resulting in ground-state restructuring, which in the large-particle limit is manifested by nonanalyticity in the ground-state energy of the quantum Hamiltonian [Eq. (4)].

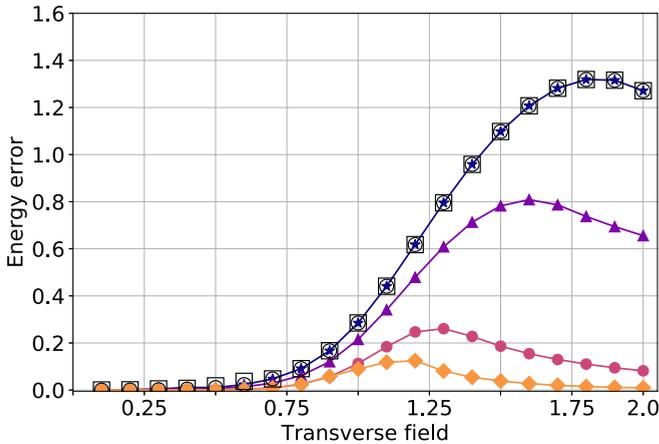


FIG. 5. Absolute value of the difference in energy between the exact solution and VQE solutions for the transverse-field Ising model. Hollow squares: rank-1 ansatz; hollow circles: tree tensor network; filled markers: checkerboard states (\star : 1 layer; \blacktriangle : 2 layers; \bullet : 3 layers; \blacklozenge : 4 layers).

It is therefore intuitively clear that, in the absence of magnetic field, $h = 0$, or in the case of high spin polarization, $h = \infty$, the Hamiltonian is dominated by a single interaction, making the ground state disentangled so that even rank-one approximations could provide quantitatively correct results. Meanwhile, this is not the case at criticality, $h = J$, where the competition between the two mechanisms favors the formation of highly entangled ground state(s).

In our numerical implementation, we use QISKIT [35] to simulate quantum circuits and the limited Broyden-Fletcher-Goldfarb-Shanno method (L-BFGS-B) to update the parameters during the classical step of VQE. We scan values of h from 0 to $2J$. For $h = 0$, the optimization process started from a random point, then each additional point begins from the previous solution. To eliminate any obviously suboptimal solutions, we also run the scanning in the opposite direction and for each value of the field we keep the better result. Without this “double-sweeping” procedure, spurious solutions appear in the vicinity of the

phase transition: the solution for one phase remains a local optimum for some time after the parameter has moved to the other phase (see Appendix C).

All ansatz states show an increase in error near the phase-transition point (Fig. 5). With the increasing depth, the checkerboard states show a better approximation. At $h = 0$, the ground state is degenerate. At nonzero values of the field, this degeneracy is lifted, but the spectral gap is very small. Nonetheless, the overlap of the VQE solution in this regime is never smaller than $1/2$ and gradually moves to 1 across the phase transition. At $h = 1.2$, the point where the energy error is the worst, the overlap of the VQE solution with the ground state is equal to 0.92. The low-depth ansatz states, namely, rank-one, tree tensor network, and checkerboard states of depth one, all show similar results, which is associated with the relative simplicity of the Ising Hamiltonian. To feed the dataset into the classifier we prepare 100 data points using VQE with four-layered checkerboard ansatz states. These data were then shuffled and split into the training set (80%) and the test set (20%). The accuracy of the prediction achieves 99%. The outcome of the quantum classifier is presented in Fig. 6 (left).

Another exactly solved model which we use to test our classifier is the antiferromagnetic XXZ spin chain with the Hamiltonian

$$H = \sum_{i=1}^n [J_{\perp}(\sigma_i^x \sigma_{i+1}^x + \sigma_i^y \sigma_{i+1}^y) + J_z \sigma_i^z \sigma_{i+1}^z]. \quad (5)$$

From a physical perspective, Eq. (5) corresponds to a uniform exchange-coupled system with a uniaxial anisotropy specified by J_z . At $|J_z| < J_{\perp}$, this model is in the XY , or planar, phase, which is characterized by algebraic decay of equal-time spin-spin correlation functions. In the regime $J_z > J_{\perp}$ the Hamiltonian corresponds to the antiferromagnetic Ising state. The system undergoes a Berezinsky–Kosterlitz–Thouless type phase transition at $J_z = J_{\perp}$ [36]. At the phase-transition point, the ground state has the highest nearest-neighbor concurrence and a cusp in nearest-neighbor quantum discord [37].

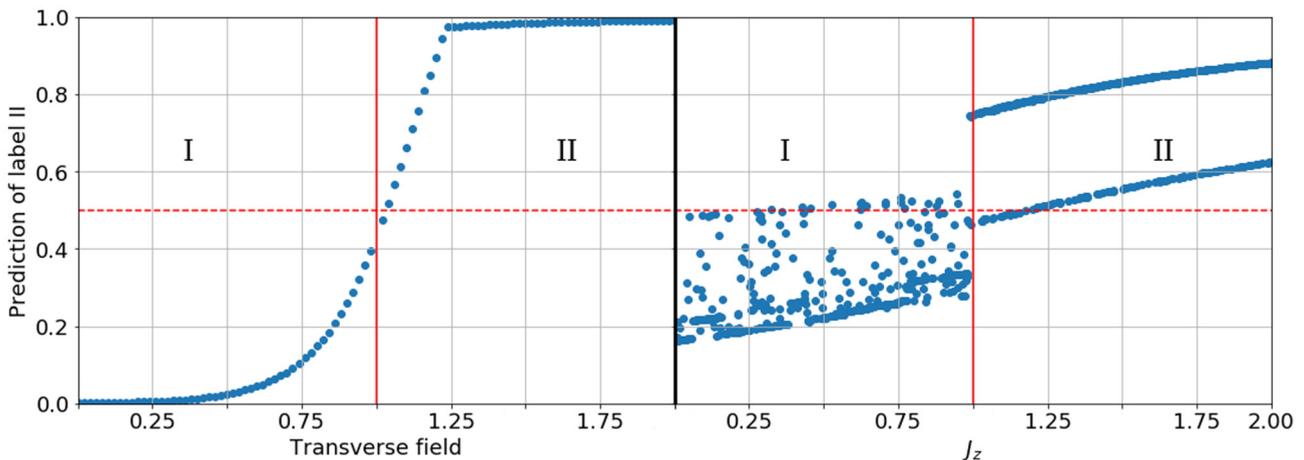


FIG. 6. (left) Predicted label of phase II as a function of magnetic field for transverse-field Ising model. (right) Predicted label of phase II as a function of J_z for the XXZ model. Roman numbers denote the phases I and II of the models.

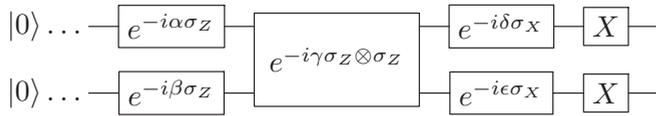


FIG. 7. Two-qubit entangler gate with additional X rotations applied.

This model has symmetry with respect to rotations in the xy plane, as well as spin-flip symmetry. This fact allows us to augment the training data. Given a VQE approximation, we can create another, equally valid approximation by applying a rotation or spin flip. The structure of the VQE ansatz is conserved in the sense that the new states are produced by the same quantum circuit with different control parameters (see Appendix A). In total, we produce 4000 data points.

The phase transition in the XXZ model is correctly learned by the classifier, yielding correct labels on 94.6% of test data (Fig. 6, right). In this case, we added two more layers to the classifier circuit to increase the accuracy. Notably, the plot for the XXZ model looks less uniform than the plot for the Ising model. This is partially connected to the data-augmentation procedure described in Appendix A.

VI. CONCLUSION

We proposed a method of classifying phases of matter that works intrinsically with quantum data, providing an advantage over the classical methods based on Monte Carlo sampling. It is a nontrivial fact that the Ising model requires fewer layers than the XXZ model. In the transverse-field Ising model, the magnetization $\sum \langle \sigma_x^{(i)} \rangle$ as a function of magnetic field clearly points at the location of the phase-transition points. This implies that the phases of the model are easy to classify. In the XXZ model, the transition at $J_z = 1$ is a transition between a paramagnetic and an antiferromagnetic phase [36]. Neither of these phases shows spontaneous magnetic moment in absence of an external field, making it somewhat harder to discern the two phases. The proposed classification technique can be applied to any model that can be expressed as a spin model (e.g., fermion problems can be mapped to spin problems by using a Jordan–Wigner transformation or a Bravyi–Kitaev transformation).

VII. NUMERICAL METHODS

The quantum circuits designed in the study were simulated by using Qiskit [35]. Figures 5, 6, 8, and 9 were prepared by using MATPLOTLIB [38]. The source code for the study is available at [39].

ACKNOWLEDGMENTS

We thank D. Yudin for fruitful discussions and suggestions. A.U. acknowledges RFBR under Project No. 19-31-90159.

APPENDIX A: DATA AUGMENTATION FOR XXZ MODEL

The XXZ Hamiltonian is symmetric with respect to spin flips and xy plane rotations. Despite the fact that its ground

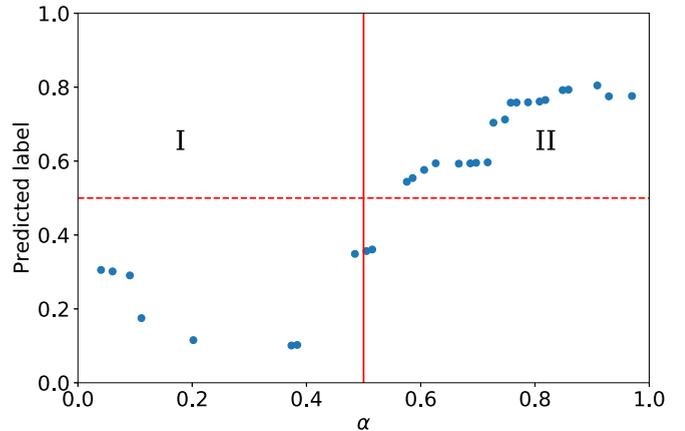


FIG. 8. Results of the learning on the random Hamiltonians model.

state is nondegenerate, applying these symmetries to a VQE state produces a different state with the same energy, which is an equally valid data point. Thankfully, these actions can be easily performed on the checkerboard states without changing their structure.

Let us start by considering rotation symmetry. This rotation is implemented by applying a Z rotation to each qubit:

$$U_{\text{rot}} = (e^{i\varphi Z})^{\otimes n}. \quad (\text{A1})$$

In the ansatz we developed, the two-qubit blocks precede Z rotations. So, applying this symmetry amounts to changing the angles in the Z rotations of the last checkerboard layer by φ .

Somewhat more complicated is the application of spin flips. We consider spin flips as applying one of X , Y , or Z operations to all spins simultaneously. The Z spin flip is a special case of the Z rotation(s). The Y flip can be composed out of X and Z flips, so we only need consider the X flip. Consider the quantum circuit in Fig. 7.

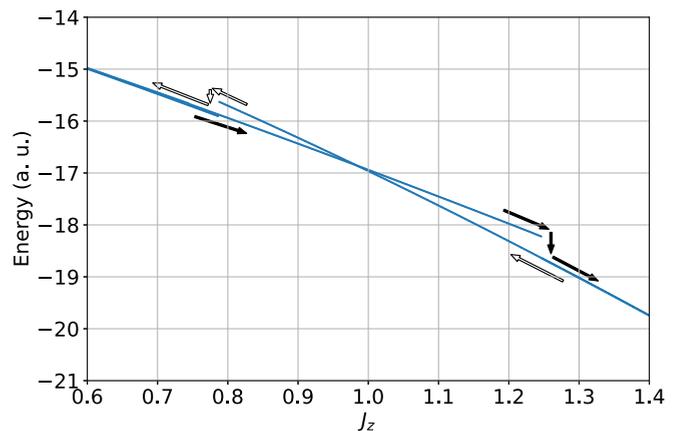


FIG. 9. Ground-state energy estimate for the XXZ model found in VQE sweeps. Filled (empty) arrows guide the eye along the “up” (“down”) sweep. The best solution out of two sweeps was subsequently used to train the classifier.

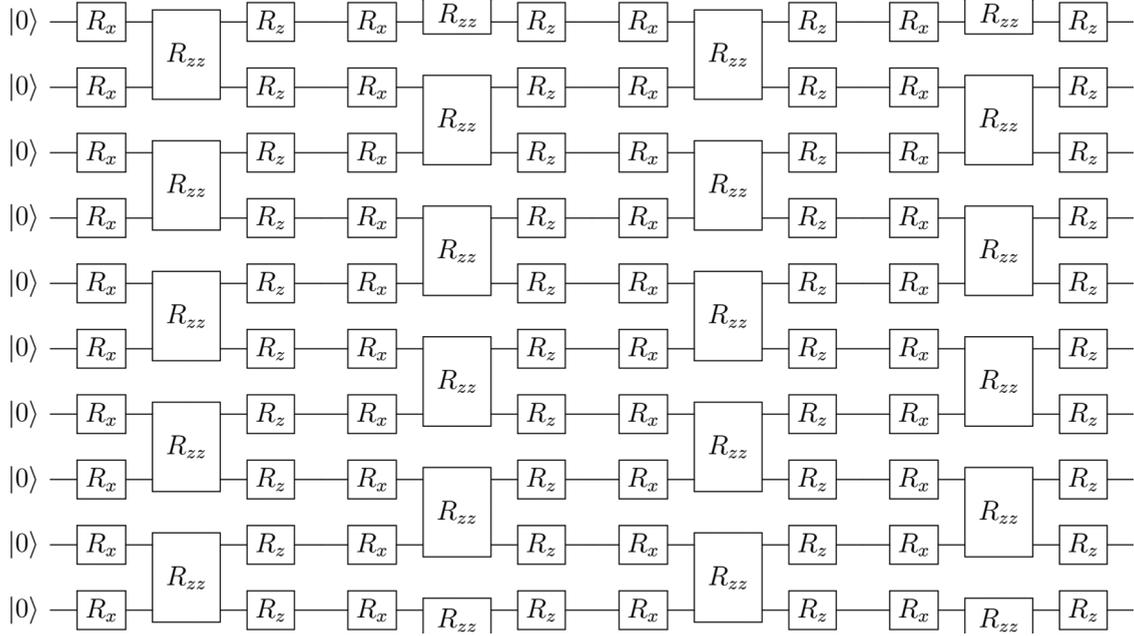


FIG. 10. Seventy-five-parameter ten-qubit circuit representing the classifier U_{class} . Here $R_\sigma \equiv R_\sigma(\theta) = e^{-i\theta\sigma}$ and $R_{\sigma\sigma} \equiv R_{\sigma\sigma}(\theta) = e^{-i\theta\sigma\otimes\sigma}$, $\sigma \in \{\mathbb{1}, X, Y, Z\}$.

Let us use the fact that the X and Z Pauli matrices anticommute and push the X gates to the front:

$$\begin{aligned} X e^{i\theta Z} &= \cos \theta X + i \sin \theta XZ \\ &= (\cos \theta \mathbb{1} - i \sin \theta Z)X = e^{-i\theta Z}X. \end{aligned} \quad (\text{A2})$$

Thus, if we push the X gates to the left and invert the angles of the Z rotations, the circuit remains invariant. Now, the next gate is the $Z \otimes Z$ rotation. $[X \otimes X, Z \otimes Z] = 0$; therefore, the X gates can go through the R_{ZZ} gate without any changes. Finally, the X gates merge with the X rotations by incrementing the angle by $\pi/4$. Thus, to augment the data with the spin-flipped states, one inverts the angles of the Z rotations in the last layer and increases the angles of the X rotations in the last layer by $\pi/4$.

APPENDIX B: TESTING THE PROCESSOR ON A MODEL WITHOUT STRUCTURE

As pointed out in the main text, simpler toy models may have simple classification criteria which do not require the application of machine learning. In this section, we classify the solutions of a randomized model: $H(\alpha) = (1 - \alpha)H_1 + \alpha H_2$, $\alpha \in [0, 1]$, where H_1 and H_2 are random Hermitian matrices pulled from a Gaussian unitary ensemble. We split the solutions into two classes: (i) $\alpha < 0.5$ and (ii) $\alpha > 0.5$. Then, we run the optimization routine to train the learning circuit to discern between the two classes.

The approach was tested for six qubits, $n = 100$, where $n_{\text{train}} = 70$, $n_{\text{test}} = 30$. The depth of the VQE circuit and the classifier circuit were both set to four layers.

The results are shown in Fig. 8. For this configuration, the accuracy of 93% was reached. This shows that the algorithm works even for such a low-structured problem, although the factors affecting the performance still require further investigation.

APPENDIX C: ON CONVERGENCE OF VARIATIONAL QUANTUM EIGENSOLVER WITH WARM STARTS

To save time on VQE computations we used the previously found solution for $H(h)$ as a starting point for the VQE process on the next Hamiltonian $H(h + \Delta_h)$ (this approach is also known as adiabatic-assisted VQE [41]). Since the Hamiltonian is deformed only slightly, the previous point is a good guess for the new minimum. Unfortunately, if during the deformation of the Hamiltonian this local minimum stops being a global one, the solver gets stuck in the wrong solution for a while. This is exactly what happens in the vicinity of the phase transition. Figure 9 demonstrates the behavior of the VQE solution for the XXZ model. We ran VQE in two sweeps: J_z swept from 0 to 2 in the “up” sweep and from 2 to 0 in the “down” sweep. As a result, VQE shows a hysteresis.

APPENDIX D: A DETAILED DEPICTION OF THE QUANTUM CIRCUIT

Figure 10 shows the classifier circuit (four layers) in full detail for ten qubits. A full circuit (VQE + classifier) can be obtained by concatenating two copies of that circuit.

- [1] J. Carrasquilla and R. G. Melko, *Nat. Phys.* **13**, 431 (2017).
- [2] E. van Nieuwenburg, Y.-H. Liu, and S. Huber, *Nat. Phys.* **13**, 435 (2017).
- [3] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature (London)* **549**, 195 (2017).
- [4] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, *Quantum Sci. Technol.* **4**, 024001 (2019).
- [5] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, *Phys. Rev. A* **101**, 032308 (2020).
- [6] M. Schuld and N. Killoran, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [7] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, *Nature (London)* **567**, 209 (2019).
- [8] M. Schuld, I. Sinayskiy, and F. Petruccione, *Contemp. Phys.* **56**, 172 (2015).
- [9] B. Duan, J. Yuan, Y. Liu, and D. Li, *Phys. Rev. A* **96**, 032301 (2017).
- [10] B. Duan, J. Yuan, J. Xu, and D. Li, *Phys. Rev. A* **99**, 032311 (2019).
- [11] Y.-B. Sheng and L. Zhou, *Sci. Bull.* **62**, 1025 (2017).
- [12] I. M. Georgescu, S. Ashhab, and F. Nori, *Rev. Mod. Phys.* **86**, 153 (2014).
- [13] H. Bernien, S. Schwartz, A. Keesling, H. Levine, A. Omran, H. Pichler, S. Choi, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin, *Nature (London)* **551**, 579 (2017).
- [14] J. T. Barreiro, M. Müller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt, *Nature (London)* **470**, 486 (2011).
- [15] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O'Brien, *Nat. Commun.* **5**, 4213 (2014).
- [16] L. Wang, *Phys. Rev. B* **94**, 195105 (2016).
- [17] J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik, *New J. Phys.* **18**, 023023 (2016).
- [18] D. Wecker, M. B. Hastings, and M. Troyer, *Phys. Rev. A* **92**, 042303 (2015).
- [19] V. Akshay, H. Philathong, M. E. S. Morales, and J. D. Biamonte, *Phys. Rev. Lett.* **124**, 090504 (2020).
- [20] P. J. J. O'Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, A. Megrant, J. Y. Mutus, C. Neill *et al.*, *Phys. Rev. X* **6**, 031007 (2016).
- [21] J. I. Colless, V. V. Ramasesh, D. Dahlen, M. Blok, M. E. Kimchi-Schwartz, J. R. McClean, J. Carter, W. A. de Jong, and I. Siddiqi, *Phys. Rev. X* **8**, 011021 (2018).
- [22] Y. Shen, X. Zhang, S. Zhang, J.-N. Zhang, M.-H. Yung, and K. Kim, *Phys. Rev. A* **95**, 020501(R) (2017).
- [23] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, *Phys. Rev. X* **8**, 031022 (2018).
- [24] E. F. Dumitrescu, A. J. McCaskey, G. Hagen, G. R. Jansen, T. D. Morris, T. Papenbrock, R. C. Pooser, D. J. Dean, and P. Lougovski, *Phys. Rev. Lett.* **120**, 210501 (2018).
- [25] O. Götze, D. J. J. Farnell, R. F. Bishop, P. H. Y. Li, and J. Richter, *Phys. Rev. B* **84**, 224428 (2011).
- [26] A. G. Taube and R. J. Bartlett, *Int. J. Quantum Chem.* **106**, 3393 (2006).
- [27] B. P. Lanyon, J. D. Whitfield, G. G. Gillet, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, A. Aspuru-Guzik, and A. G. White, *Nat. Chem.* **2**, 106 (2010).
- [28] Y. Cao, J. Romero, J. P. Olson, M. Degroote, P. D. Johnson, M. Kieferová, I. D. Kivlichan, T. Menke, B. Peropadre, N. P. D. Sawaya, S. Sim, L. Veis, and A. Aspuru-Guzik, *Chem. Rev.* **119**, 10856 (2019).
- [29] F. Vatan and C. Williams, *Phys. Rev. A* **69**, 032315 (2004).
- [30] R. Santagati, J. Wang, A. A. Gentile, S. Paesani, N. Wiebe, J. R. McClean, S. Morley-Short, P. J. Shadbolt, D. Bonneau, J. W. Silverstone, D. P. Tew, X. Zhou, J. L. O'Brien, and M. G. Thompson, *Sci. Adv.* **4**, eaap9646 (2018).
- [31] E. Farhi, J. Goldstone, and S. Gutmann, [arXiv:1411.4028](https://arxiv.org/abs/1411.4028).
- [32] J. Spall, *IEEE Trans. Autom. Control* **37**, 332 (1992).
- [33] E. Lieb, T. Schultz, and D. Mattis, *Ann. Phys. (NY)* **16**, 407 (1961).
- [34] P. Pfeuty, *Ann. Phys. (NY)* **57**, 79 (1970).
- [35] G. Aleksandrowicz *et al.*, Qiskit: An Open-Source Framework for Quantum Computing (2019).
- [36] F. Franchini, *An Introduction to Integrable Techniques for One-Dimensional Quantum Systems*, Lecture Notes in Physics (Springer International Publishing, Cham, 2017), Vol. 940.
- [37] R. Dillenschneider, *Phys. Rev. B* **78**, 224413 (2008).
- [38] J. D. Hunter, *Comput. Sci. Eng.* **9**, 90 (2007).
- [39] github.com/Quantum-Machine-Learning-Initiative/learning-phases.