


Tensor-network discriminator architecture for classification of quantum data on quantum computers

Michael L. Wall ,* Paraj Titum, and Gregory Quiroz

The Johns Hopkins University Applied Physics Laboratory, Laurel, Maryland 20723, USA

Michael Foss-Feig

Quantinuum, 303 S. Technology Ct., Broomfield, Colorado 80021, USA

Kaden R. A. Hazzard

*Department of Physics and Astronomy, Rice University, Houston, Texas 77005-1892, USA
and Rice Center for Quantum Materials, Rice University, Houston, Texas 77005-1892, USA*



(Received 24 February 2022; accepted 7 June 2022; published 22 June 2022)

We demonstrate the use of matrix product state (MPS) models for discriminating quantum data on quantum computers using holographic algorithms, focusing on the problem of classifying a translationally invariant quantum state based on L qubits of quantum data extracted from it. We detail a process in which data from single-shot experimental measurements are used to optimize an isometric tensor network, the isometric tensors are compiled into unitary quantum operations using greedy compilation heuristics, parameter optimization on the resulting quantum circuit model removes the postselection requirements of the isometric tensor model, and the resulting quantum model is inferred on either product state (single-shot measurement) or entangled quantum data. We demonstrate our training and inference architecture on a synthetic dataset of six-site single-shot measurements from the bulk of a one-dimensional transverse field Ising model (TFIM) deep in its antiferromagnetic and paramagnetic phases. We find that increasing the bond dimension of the tensor-network model, amounting to adding more ancilla qubits to the circuit representation, improves both the average number of correct classifications across the dataset and the single-shot probability of correct classification. We experimentally evaluate models on Quantinuum's H1-2 trapped ion quantum computer using entangled input data modeled as translationally invariant, bond dimension 4 MPSs across the known quantum phase transition of the TFIM. Using linear regression on the experimental data near the transition point, we find predictions for the critical transverse field of $h = 0.962$ and 0.994 for tensor-network discriminators of bond dimension $\chi = 2$ and $\chi = 4$, respectively. These predictions compare favorably with the known transition location of $h = 1$ despite training on data far from the transition point. Our techniques identify families of short-depth variational quantum circuits in a data-driven and hardware-aware fashion and robust classical techniques to precondition the model parameters, and can be adapted beyond machine learning to myriad applications of tensor networks on quantum computers, such as quantum simulation and error correction.

DOI: [10.1103/PhysRevA.105.062439](https://doi.org/10.1103/PhysRevA.105.062439)

I. INTRODUCTION

Quantum technology continues to advance rapidly. This progress urgently requires developing and assessing applications of near-term quantum technology and methods to analyze and interpret the large volumes of quantum information these devices will produce. Beyond-classical computational tasks [1] have been demonstrated, including sampling from random unitaries [2] and boson sampling [3], and smaller-scale demonstrations and proposals exist for quantum simulation [4], optimization [5], machine learning [6], and other topics [7]. In contrast to the widely used randomized benchmarking [8] and random unitary simulation [2] use cases, applications like quantum simulation and machine

learning will deal with physical data that is expected to be highly structured. Identifying quantum models that can discover and adapt to this underlying structure, as is the case in successful classical approaches to similar problems [9,10], is key to both developing performant algorithms and identifying hardware improvements that improve model performance in relevant metrics.

Tensor networks (TNs), in which a high-rank tensor is decomposed into a contracted network of low-rank tensors, have been used in a variety of applications across physics, chemistry, applied mathematics, and machine learning [11–17]. In quantum physics, TNs are used to represent the coefficients of a wave function in a fixed basis, and the geometry of the TN is related to the entanglement and correlation structure of the resulting wave function [18,19]. Variational methods within the class of TNs define the gold standard for simulations of strongly correlated quantum many-body systems on classical

*Michael.Wall@jhuapl.edu

computers, with the most prominent example being the density matrix renormalization group (DMRG) algorithm [11,20]. Similar to how classical machine learning has benefited significantly by tailoring the structure of the learning model to the structure of the data, e.g., to account for its hierarchical or causal nature [21], we can expect that models which encode quantum characteristics of the data may be more performant or interpretable when applied to quantum learning problems.

There is also strong impetus for developing TN models to be employed on quantum computers. Certain TN architectures naturally define a sequential preparation scheme for quantum states in which qubits holding the quantum state information are correlated using a register of ancilla qubits through a sequence of quantum gates [22,23]. The ancilla qubit register is the quantum analog of the contracted indices of the classical TN representation. However, its effective bond dimension grows exponentially with the number of qubits in the register, leading to a vastly improved scaling of network expressiveness with resources compared to the classical case. The state of the ancilla qubits in a partially contracted network also carries information about the state, but scales only with the boundary of the contracted region. This so-called “holographic” encoding of quantum information [24–28] has been utilized in recent works to measure static and dynamic properties of quantum states on quantum computers [29–32]. The ability of tensor networks to partially localize information about a quantum state into a reduced-dimensional representation may also bode well for trainability in machine learning applications by avoiding the phenomenon of barren plateaus [33].

In this work, we explore the use of matrix product states (MPSs), one-dimensional tensor networks, for inferencing from quantum data on quantum computers. Namely, we focus on the problem of classifying a translationally invariant quantum state given L qubits of data extracted from it. Our discriminator architecture first generates a prior distribution on the ancilla degrees of freedom in the network, and then conditions the state of the ancilla distribution on the quantum data to be classified. A final quantum operation on the ancilla qubits encodes the class label into a subset of those qubits, which can then be read out by measurement in the computational basis. We detail a process in which the tensors in the network architecture are classically optimized using optimization techniques on manifolds inspired by canonical TN algorithms, the tensors of the optimized classical model are compiled into quantum operations using greedy compilation heuristics [34,35], and then the resulting model parameters can be fine-tuned based on results obtained from inferencing the quantum model. In contrast to many other variational quantum circuit approaches, in which a fixed-depth sequence of native gates is optimized over its parameters, our approach only defines the geometry of quantum operations between qubits and allows for an automatic determination of the circuit structure and depth by interfacing with quantum data.

We demonstrate the training of our architecture on a collection of single-shot experimental measurements, in which each shot corresponds to the measurement of L qubits from a fixed set of bases. Such single-shot measurements can be realized not only in programmable quantum hardware [36], but also in other platforms such as cold atoms [37] and trapped

ions [38,39], and so our training strategy avoids the need for the quantum system providing the data to be programmable. While the collection of single-shot measurements can be represented as vectors in a $2L$ -dimensional space and classically analyzed, the quantum model can capture that these vectors are sparse representations of vectors embedded in a 2^L -dimensional Hilbert space, and use their correlations to define entangled decision boundaries. In addition, once trained our model can be inferenced on either single-shot measurements or entangled input data, with the latter having no classical analog. We analyze the performance of both modalities in this work.

This work is organized as follows: Sec. II overviews matrix product states, applications to machine learning, and their implementation on quantum hardware; Sec. III presents our tensor-network discriminator architecture and algorithms for classical training, quantum compilation, and other classical preconditioning strategies of the quantum model; Sec. IV describes an application of our model workflow to the one-dimensional (1D) transverse-field Ising model, including classical simulations of model performance and analysis of experiments on Quantinuum’s model H1-2 trapped ion quantum computer; finally, Sec. V presents conclusions. Technical details of algorithms used to generate the synthetic training data are presented in the Appendixes.

II. MATRIX PRODUCT STATES

A. Overview

Matrix product states (MPSs), also known as finitely correlated states, are tensor networks with a 1D geometry. MPSs are of interest because they exactly represent ground states of gapped, local, 1D Hamiltonians [18,40], and also because many efficient algorithms have been developed for finding approximate MPS representations of many-body systems, such as the density-matrix renormalization group (DMRG) algorithm [20]. The MPS description of a quantum state $|\psi\rangle$ on a lattice consisting of Ld -level subsystems is given by the contraction of the L tensors $\{A^{[j]}\}$ as

$$|\psi\rangle = \sum_{i_0 \dots i_{L-1}} \langle l | \mathbb{A}^{[0]i_0} \dots \mathbb{A}^{[L-1]i_{L-1}} | r \rangle | i_0 \dots i_{L-1} \rangle, \quad (1)$$

in which $\mathbb{A}^{[j]i_j}$ denotes the matrix obtained by holding the index i_j fixed in the tensor $A^{[j]}$. Each tensor $A^{[j]}$ is (χ, d, χ) -dimensional, and the *bond dimension* χ forms the main convergence parameter of an MPS. The states $|l\rangle$ and $|r\rangle$ are χ -dimensional states specifying boundary conditions.

The matrix product structure of Eq. (1) implies that we can place any full-rank $\chi \times \chi$ matrix \mathbb{X} and its inverse \mathbb{X}^{-1} between any two neighboring tensors in the product without changing the nature of the state. This leads to a gauge freedom in MPSs that can be exploited to put the tensors into one of a number of so-called canonical forms [11]. In this work, a central role will be played by MPSs in the left-canonical form, in which $\sum_{\alpha i} A_{\alpha\beta}^{[j]i} A_{\alpha\beta'}^{[j]i*} = \delta_{\beta\beta'}$. Reshaping this tensor into a $(\chi d) \times \chi$ -dimensional matrix, left-canonical form specifies that this is an *isometric matrix* with orthonormal columns. Hence, this isometric operation can be embedded into a

unitary operation $\hat{U}^{[j]}$ such that

$$A_{\alpha\beta}^{[j]i_j} = \langle i_j | \alpha | \hat{U}^{[j]} | 0 \beta \rangle, \quad (2)$$

and this sequence of unitary operations can then be used to sequentially generate the state as

$$|\psi\rangle = \langle l | {}_b \hat{U}^{[0]} | 0 \rangle_0 \otimes \cdots \otimes \hat{U}^{[L-1]} | 0 \rangle_{L-1} \otimes | r \rangle_b. \quad (3)$$

Here the subscript b denotes states in a χ -dimensional bond Hilbert space, and the other subscripts denote qubit index in a register of L qubits used to hold the generated state.

Equation (3) codifies a scheme to prepare the state $|\psi\rangle$ using $L + \log_2 \chi$ qubits: a register of L qubits plus an additional $\log_2 \chi$ qubits to encode the bond degrees of freedom. However, this scheme can be modified to require only $\log_2 \chi + 1$ qubits for implementation provided the hardware supports midcircuit measurement and reuse (MCMR) [22,41,42]. In this scheme, $\log_2 \chi$ of the qubits form the bond register, and the remaining qubit is used to prepare the physical degrees of freedom by sequentially coupling with the bond register. Then the physical qubit information is measured before the next physical qubit state is generated, and so this reduced qubit requirement scheme can be used to calculate observables which can be expressed as tensor products of local operators. This scheme can also be used to generate sample product states in a specified basis, $|i_0 \cdots i_{L-1}\rangle$ according to their probability of occurrence in the wave function, $|\langle \psi | i_0 \cdots i_{L-1} \rangle|^2$. The $\log_2 \chi$ qubits required are exponentially smaller than the number of classical bits required to implement the MPS, and relatively large system sizes can be amenable to current devices with modest numbers of qubits [22].

MPSs can also represent the thermodynamic limit of a translationally invariant system on an infinite chain of sites obtained by the $L \rightarrow \infty$ limit of Eq. (1), in which case we can capture the state using a single, translationally invariant tensor, dropping the site index superscript in brackets. This reduced description of an infinite-size state with an MPS tensor is known as an iMPS. Intuitively, far from the boundary the particular boundary vector, e.g., $|r\rangle$ used to generate the state becomes unimportant as the state of the bond degrees of freedom reaches a steady state. This can be roughly formalized by the fact that the (right) transfer operator of the left-canonical iMPS defined by

$$T_R(A; \mathbb{E}) = \sum_i \mathbb{A}^i \mathbb{E} \mathbb{A}^{i\dagger}, \quad (4)$$

in which \mathbb{E} is a $\chi \times \chi$ matrix, admits the half-infinite reduced density matrix $\hat{\rho}_{\text{half-infinite}}$ obtained by tracing out one half of the infinite chain, as an eigenmatrix with maximal eigenvalue, which is 1 for a normalized iMPS. Hence, the half-infinite density matrix in bond space is obtained as the steady state of the channel

$$\hat{\rho} = \text{Tr}_{N_b} \hat{U} | 0 \rangle_{N_b} \otimes \cdots \otimes \text{Tr}_2 \hat{U} | 0 \rangle_2 \otimes \text{Tr}_1 \hat{U} | 0 \rangle_1 \otimes | r \rangle_b, \quad (5)$$

as $N_b \rightarrow \infty$. Here \hat{U} is the unitary embedding of the MPS tensor A in left-canonical form, defined as in Eq. (2), which acts on the combined register of the bond qubits (denoted by subscript b on the key) and a physical register (denoted by subscript numerals) before the latter is traced out. In general, tensor networks encode a D dimensional quantum system

into a register of qubits on the $(D - 1)$ -dimensional boundary, so this approach is referred to as ‘‘holographic’’ [24–29]. Holographic approaches have been recently exploited in a trapped-ion quantum computer to measure the entanglement entropy defined by the spectrum of $\hat{\rho}_{\text{half-infinite}}$ [30,31], as well as for dynamics [32]. In practice, since MPSs of finite bond dimension have a finite correlation length, a finite ‘‘burn-in length’’ N_b suffices (set to L in the results that follow), and the boundary state $|r\rangle$ can be chosen to facilitate rapid convergence to the steady state (see Appendix A for further algorithmic details).

B. MPSs for machine learning

Given that MPSs define a sequential preparation procedure for quantum states starting from a given reference state, a natural application of MPSs is generative modeling [34,43]. In this scenario, the tensors A of an MPS are optimized to best represent a collection of training quantum states $\{|\psi_m\rangle\}$, $m = 1, \dots, M$ by some metric, such as the average log-likelihood of the data in the MPS

$$\mathcal{L} = \frac{1}{M} \sum_m \ln |\langle \psi_m | \psi[A] \rangle|^2, \quad (6)$$

where $|\psi[A]\rangle$ denotes the state generated by MPS tensors A . Once trained, the model can be compiled to quantum hardware by finding appropriate embeddings of the tensors A in left canonical form into unitaries as in Eq. (2) and then applying the sequential preparation procedure. A proof of principle of this generative model pipeline has been explored in Ref. [34], in which greedy heuristics were applied to compile the isometries into unitaries for hardware with a given gate set and geometry using as short of gate sequences as possible. For a finite-size system, $|r\rangle$ can be chosen to be the vacuum state of the qubits encoding the bond degrees of freedom, while for an infinite-size state the ‘‘burn-in’’ procedure described above can be enacted to first generate the half-infinite density matrix on the bond states, and then the sequential preparation of a finite portion of the state performed.

If we instead consider a supervised discriminative task, we now wish to classify a quantum state $|\psi_{\text{test}}\rangle$ into one of N_C classes given a collection of training states and class labels $\mathcal{T} = \{(|\mathbf{x}_m\rangle, \ell_m), m = 1, \dots, M\}$. One means of doing so is to define an MPS linear model [44] by introducing a $\chi \times N_C \times \chi$ -dimensional label tensor \mathbb{P}^ℓ in between the bonds of the MPS tensors in the representation Eq. (1) to define a classifier state $|\psi_{\text{class};\ell}\rangle$ as

$$|\psi_{\text{class};\ell}\rangle = \sum_{i_0 \cdots i_{L-1}} \langle l | \mathbb{A}^{[0]i_0} \cdots \mathbb{A}^{[L-1]i_{L-1}} \mathbb{P}^\ell | r \rangle | i_0 \cdots i_{L-1} \rangle. \quad (7)$$

This object is shown graphically in Fig. 1(a). The overlap of this classifier MPS with a test quantum state $|\psi_{\text{test}}\rangle$ defines an indicator function $f_\ell(|\psi_{\text{test}}\rangle)$ as

$$\begin{aligned} f_\ell(|\psi_{\text{test}}\rangle) &= \langle \psi_{\text{class};\ell} | \psi_{\text{test}} \rangle \\ &= \sum_{i_0 \cdots i_{L-1}} (\langle l | \mathbb{A}^{[0]i_0} \cdots \mathbb{A}^{[L-1]i_{L-1}} \mathbb{P}^\ell | r \rangle)^* \\ &\quad \times \langle i_0 \cdots i_{L-1} | \psi_{\text{test}} \rangle, \end{aligned} \quad (8)$$

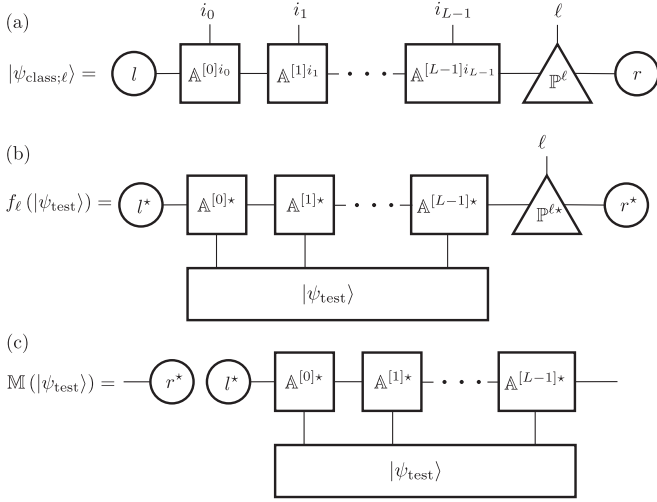


FIG. 1. MPS linear classifier model represented by tensor-network graphical notation. (a) A linear model MPS classifier is obtained by introducing a label tensor \mathbb{P}^ℓ on one of the bonds of an MPS structure. (b) The overlap of the MPS classifier wave function defined in (a) with a test wave function $|\psi_{\text{test}}\rangle$ defines an indicator function $f_\ell(|\psi_{\text{test}}\rangle)$ that can be optimized for classification performance. (c) The action of the linear MPS model can be interpreted as an encoding of the test state information into the bond space to define an operator $\mathbb{M}(|\psi_{\text{test}}\rangle)$ shown graphically, followed by classification using the indicator tensor as $\text{argmax}_\ell \text{Tr}(\mathbb{P}^{\ell\dagger} \mathbb{M}(|\psi_{\text{test}}\rangle))$.

shown graphically in Fig. 1(b). This indicator function can then be optimized over the tensors of the MPS (A and P) such that $f_\ell(|\psi_{\text{test}}\rangle) \approx 1$ when $|\psi_{\text{test}}\rangle$ is in class ℓ and $f_{\ell'}(|\psi_{\text{test}}\rangle) \approx 0$ for all other classes ℓ' , with performance averaged over the training set \mathcal{T} . We note that the position of the indicator tensor \mathbb{P}^ℓ in the MPS chain can be shifted using gauge freedom, and so we have shifted it to the right end for convenience. We can rewrite the indicator function as $f_\ell(|\psi_{\text{test}}\rangle) = \text{Tr}(\mathbb{P}^{\ell\dagger} \mathbb{M}(|\psi_{\text{test}}\rangle))$ in which

$$\mathbb{M}(|\psi_{\text{test}}\rangle) = \sum_{i_0 \cdots i_{L-1}} A^{[L-1]i_{L-1}\dagger} \cdots \times A^{[0]i_0\dagger} \langle i_0 \cdots i_{L-1} | \psi_{\text{test}} \rangle |l\rangle \langle r|. \quad (9)$$

$\mathbb{M}(|\psi\rangle)$ is shown graphically in Fig. 1(c) and represents the overlap between the vector $|\psi\rangle$ and the state represented by the MPS tensors A in bond space. Hence, the indicator tensor \mathbb{P}^ℓ performs classification based on the bond space encoding of this overlap. In this way, an MPS classifier performs a lower-dimensional ‘‘holographic’’ encoding of the information needed for classification into the bond degrees of freedom, in analogy to how classical encoders produce reduced dimensional representations in an appropriately defined latent space.

This observation leads to a duality between generative and discriminative MPS models. Namely, consider that we have two generative MPS models with tensors A and A' and bond dimensions χ and χ' encoding training sets \mathcal{T} and \mathcal{T}' with high fidelity. If we now consider that the elements of \mathcal{T} are

in class 0 and those of \mathcal{T}' are in class 1, we can construct an MPS classifier with bond dimension $\chi_c = \chi + \chi'$ as

$$\begin{aligned} |\psi_{\text{class};\ell}\rangle &= \mathcal{N} \sum_{i_0 \cdots i_{L-1}} |i_0 \cdots i_{L-1}\rangle \langle l| \otimes \langle l'| \\ &\times (A^{[0]i_0} \otimes \mathbb{I}_{\chi'} + \mathbb{I}_\chi \otimes A'^{[0]i_0}) \\ &\times \cdots (A^{[L-1]i_{L-1}} \otimes \mathbb{I}_{\chi'} + \mathbb{I}_\chi \otimes A'^{[L-1]i_{L-1}}) \\ &\times (\delta_{\ell,0} \hat{P}_0 + \delta_{\ell,1} \hat{P}_1) |r\rangle \otimes |r'\rangle, \end{aligned} \quad (10)$$

in which $\delta_{\ell,\ell'}$ is the Kronecker delta, \hat{P}_0 (\hat{P}_1) is a projector onto the bond states connecting to the tensors A (A'), and \mathcal{N} is a normalization factor. The indicator function associated with this MPS model applied to a test state $|\mathbf{x}\rangle$ is

$$f_\ell(|\mathbf{x}\rangle) = \mathcal{N} (\delta_{\ell,0} \langle \psi[A] | \mathbf{x} \rangle + \delta_{\ell,1} \langle \psi[A'] | \mathbf{x} \rangle). \quad (11)$$

Hence, this indicator function behaves as desired, selecting the class according to the state with higher overlap, and will be accurate so long as the two classes are sufficiently distinguishable (have small overlap $\ll 1$). In the extreme case that one of the generative models is provided as a test state, the correct class will have weight $\propto 1$ while the incorrect class will be weighted by a factor $\propto \langle \psi[A] | \psi[A'] \rangle$ that will be $\ll 1$ when the classes are distinguishable.

The above process generates an accurate classifier provided that the training data within each class can be well described as an MPS with the specified bond dimension. However, there are several drawbacks of this approach. Perhaps most strikingly, this process requires building of generative models which capture all properties of the training data, and then requires all of this information from all classes to be encoded into the classifier model. This leads to potentially very large bond dimensions. However, just as in classical learning, it is likely that the classification decision can be made based on a set of features derived from the data, whose size is much smaller than the size of the full data set. In the architecture we present in the next section, we avoid this drawback by training a generator in feature space and the discriminator simultaneously. This enables us to model only the parts of the data necessary for the classification decision, reducing the resource requirements for the model. Further, this approach does not require that the training data have an efficient MPS representation, but only that the set of features used for the classification decision does.

The second point we will address with our tensor-network discriminator architecture has to do with the implementation of a classifier like Eq. (10) on quantum hardware. As follows from Eqs. (10) and (3), an MPS classifier can be applied to a separable feature state $|\mathbf{x}\rangle = |x_0\rangle \otimes \cdots \otimes |x_{L-1}\rangle$ using $\log_2 \chi + \lceil \log_2 N_c \rceil$ qubits and MCMR as

$$\langle r | {}_b \hat{\mathcal{U}} | 0 \rangle_c \langle 0 | \hat{\mathcal{U}}^{[L-1]\dagger} | x_{L-1} \rangle \otimes \cdots \otimes \langle 0 | \hat{\mathcal{U}}^{[0]\dagger} | x_0 \rangle \otimes |l\rangle_b, \quad (12)$$

in which $\hat{\mathcal{U}}^{[l]}$ is a unitary embedding of the MPS tensor defining the generative interpretation of the model Eq. (2) and $\hat{\mathcal{U}}$ is a similar unitary embedding of \mathbb{P}^ℓ such that

$$P_{\alpha\beta}^\ell = \langle \ell \alpha | \mathcal{U} | 0 \beta \rangle, \quad (13)$$

in which ℓ is the class label interpreted in binary. Following several executions of the circuit, the most probable inferred

value of ℓ obtained by measurements in the computational basis can be assigned as the predicted class label.

As is evident from Eq. (12), this circuit also requires postselection on the physical qubits being in the state $|0\rangle$ following action of $\hat{U}^{[j]}$. When the state $|1\rangle$ is instead measured it indicates that the action of the circuit does not correspond to action of the tensor $A^{[j]}$ defining the MPS, but rather to the portion of the unitary embedding $\hat{U}^{[j]}$ that is not determined uniquely by the MPS representation. While the probability of successful postselection may be high in the case that the training and test feature vectors have significant overlap and the MPS model is very expressive, it is clearly preferred to develop a model that does not require these conditions to be met. In the next section we develop an architecture inspired by these observations that will obviate the need for postselection.

III. TENSOR-NETWORK DISCRIMINATOR ARCHITECTURE

In this section, we put forth a tensor-network discriminator architecture for quantum hardware that enables efficient classical training and preconditioning, avoids the need for postselection that arises from naive application of an MPS model for classification, and produces local cost functions with a data-defined circuit architecture amenable to refinement and optimization with the quantum device in the loop. While our methodology can also be straightforwardly adapted to finite-size systems with boundaries, we will henceforth specialize to the case of learning from a finite string sampled from an infinite, translationally invariant state.

Our discriminator consists of four key model components to be trained, which we will denote as \hat{U}_R , \hat{U}_G , \hat{U}_D , and \hat{U}_C . \hat{U}_R is a $\chi \times \chi$ unitary operator acting on the bond degrees of freedom to form an appropriate representation of the right boundary condition $|r\rangle$. \hat{U}_G is a $(\chi d) \times (\chi d)$ -dimensional unitary that generates an appropriate prior distribution for the feature space in the bond degrees of freedom starting from the state generated by \hat{U}_R . Together, \hat{U}_R and \hat{U}_G generate the appropriate boundary conditions and prior information for the classifier, similar to the contribution $|l\rangle\langle r|$ in the matrix \mathbb{M} defined in Eq. (9). \hat{U}_D is a $(\chi d) \times (\chi d)$ -dimensional unitary that conditions the state of the bond degrees of freedom based on an input vector $|\mathbf{x}\rangle$. Finally, \hat{U}_C is a $(\chi d) \times (\chi d)$ -dimensional¹ unitary that encodes predicted class labels into the probabilities of the states in a $\log_2 N_C$ -dimensional register based on the state of the bond degrees of freedom output from \hat{U}_D . In this way, \hat{U}_G is analogous to the unitaries \hat{U} generating an MPS state as in Eqs. (2) and (3), \hat{U}_D is analogous to the adjoints of \hat{U} utilized in an MPS classifier as Eq. (12), and \hat{U}_C classifies based on the state of the bond degrees of freedom analogous to \hat{U} in Eqs. (12) and (13). In addition to the input parameters L (the number of qubits encoding the input data) and N_C (the number of classes) there are two hyperparameters of the model: N_b is the number of times \hat{U}_G is applied to

reach a steady-state bond register prior distribution and χ is the bond dimension.

This process is displayed visually in Fig. 2 for the case of a product input vector $|\mathbf{x}\rangle$. The total qubit requirements for the classifier, assuming MCMR, are $\log_2 \chi + 1$. This does not include the resources required to encode the input vector $|\mathbf{x}\rangle$. In the general case that this is an arbitrary quantum state it will require L qubits to encode. Two important special cases are (1) when all input vectors are product states, in which case the physical qubit used in the classifier can be utilized and no additional resources are required, and (2) when the input data are described by an MPS with bond dimension χ' , in which case an additional $\log_2 \chi'$ qubits are required. We will demonstrate both of these modalities in Sec. IV.

A. Classical training and preconditioning of tensor-network architecture

We now turn to how the tensor-network discriminator architecture can be optimized classically using adaptations of standard tensor-network techniques. This classical preconditioning serves several benefits. For one, the classical preconditioning will define a quantum circuit architecture with initial guesses that can then be optimized further on the quantum device, with the initialization helping to avoid local minima and barren plateaus [45–47]. In stark contrast to most other approaches for variational circuit learning, in which a predefined circuit architecture is optimized or operations are characterized as general dense unitaries without reference to their representation in native gates for a given hardware realization, our approach will produce a family of quantum circuits whose depth and native gate content are discovered from the correlation structure of the training data itself. Finally, from a benchmarking perspective this process produces models that can be analyzed on either classical or quantum hardware to assess the quality of noise models for the latter in structured use cases.

The operators \hat{U}_R , \hat{U}_G , \hat{U}_D , and \hat{U}_C can be interpreted as unitary embeddings of tensors R , G , D , and C , analogous to the embeddings described in Sec. II. R describes a normalized χ -dimensional state, and so is an element of the sphere manifold of dimension χ , S^χ , whose elements have unit norm,² and the unitary embedding takes the form $R_\alpha = \langle \alpha | \hat{U}_R | 0 \rangle$. The tensor G can be described as a (χ, d, χ) -dimensional tensor in left canonical form, such that the $(\chi d) \times \chi$ -dimensional matrix reshaping of this tensor is an element of the Stiefel manifold $\text{St}(\chi d, \chi)$ of matrices with orthonormal columns. Similarly, D is a (χ, d, χ) -dimensional tensor where the adjoint of the $\chi \times (\chi d)$ -dimensional matrix reshaping is an element of $\text{St}(\chi d, \chi)$. The unitary embeddings of these tensors are analogous to Eq. (2) for MPS tensors. Finally, C is a $\chi \times N_C$ dimensional matrix that is an element of $\text{St}(\chi, N_C)$ and embedded as $C_{\alpha,\ell} = \langle 0 | {}_b \langle \ell | \hat{U}_C | \alpha \rangle_b | 0 \rangle$, where we have assumed binary classification for simplicity.

Given that all of the tensors in our architecture are elements of Riemannian manifolds, using manifold optimization techniques is a natural approach. Several recent works have

¹ \hat{U}_C is not required to act on the physical qubit; however, since the physical qubit is needed for other parts of the architecture, we have included it as a potential resource to be utilized by \hat{U}_C .

²Throughout the term norm refers to the L_2 norm.




FIG. 2. Diagram of tensor-network discriminator architecture, with operation read left to right as in quantum circuit diagrams. First, the operator \hat{U}_R creates an initial state in the $\log_2 \chi$ bond qubits, starting from the vacuum state. Next, this bond state is coupled with a physical qubit via the operator \hat{U}_G and the physical qubit state discarded and reset N_b times. This creates a bond state that can be interpreted as a feature space representation of the density matrix obtained by tracing out half of the infinite system. This bond state is now combined with the input state $|\mathbf{x}\rangle$ a single qubit at a time via the operator \hat{U}_D , and the physical qubit measured and reinitialized after each application. For clarity, this operation is drawn as though the test state is a product state, but the process also applies to any entangled input state stored in an L -qubit register. Finally, the resulting bond degrees of freedom which have been conditioned on the observation of the state $|\mathbf{x}\rangle$ are operated on by \hat{U}_C , and the measurement of $\log_2 N_C$ of these qubits defines a binary representation of the predicted class label.

discussed manifold optimization techniques, including within the context of tensor networks and other quantum applications [48,49]. In order to apply these techniques, we need to define a cost function for the optimization over a training set $\mathcal{T} = \{|\mathbf{x}_m\rangle, \ell_m, m = 1, \dots, M\}$. The key output to be optimized is the density matrix ρ of the class qubit register following the process outlined in Fig. 2. Namely, we would like to optimize ρ such that the elements $\rho_{\ell_m \ell_m}$, corresponding to the probability of measuring $|\ell_m\rangle$ when inferring a state $|\mathbf{x}_m\rangle$, are as large as possible, while all other probabilities $\rho_{\bar{\ell} \bar{\ell}}$ are as small as possible, averaged over the training set. When evaluated with the tensors R, G , etc. rather than their unitary embeddings, the analog of this density matrix will no longer have unit trace and so is not a true density matrix. Hence, in order to treat all training states on the same footing, this object should be divided by its trace when evaluated in this fashion. In practice, we find that minimizing the cost function $\frac{1}{M} \sum_m \mathcal{C}(\rho^m, \ell_m)$, in which

$$\mathcal{C}(\rho, \ell_m) = \frac{\sum_{\ell} \rho_{\ell \ell} - 2\rho_{\ell_m \ell_m}}{\sum_{\ell} \rho_{\ell \ell}}, \quad (14)$$

performs well. We note that the coherences (off-diagonal elements) of the density matrix play no role in the optimization, since the final readout is in the computational basis.

Summarizing, we can define a cost function amenable to manifold optimization over the tensors R, G, D, C as in Algorithm 1, which assumes a training set of product state vectors $|\mathbf{x}_m\rangle = |x_{m,0}\rangle \cdots |x_{m,L-1}\rangle = (\sum_{i_0} x_{m,0}^{i_0} |i_0\rangle) \cdots (\sum_{i_{L-1}} x_{m,L-1}^{i_{L-1}} |i_{L-1}\rangle)$ for simplicity. As discussed further in Sec. IV, product state samples correspond to the case of single-shot measurements relevant for near-term experiments.

A tensor-network diagram for the construction of $\rho_{\ell, \ell'}$ is shown in Fig. 3(a). In practice, we achieve the best performance when we do a first round of optimization jointly across all tensors followed by rounds of optimization where a single tensor is optimized with all others held fixed. We also note that high-quality software packages for Riemannian manifold optimization are available in a variety of programming languages [50,51].

With the optimized tensors in hand, we move to compiling the embedded unitaries to quantum circuits for deployment to quantum hardware. As in past work [34,35], we utilize a greedy iterative heuristic for compilation in which

progressively deeper circuits are constructed from past iterations until the distance between the compiled unitary and the unitary embedding of the tensor over the elements where the tensor is defined is less than a defined tolerance. The method is initialized with an ansatz consisting of single-qubit rotations on each qubit, and updates to previous iterations consist of a single, elemental two-qubit entangling gate (e.g., CNOT) together with single-qubit rotations on the modified qubits. At each iteration the circuits are optimized over their parameters (e.g., the rotation angles of the single-qubit rotations) and the best performing h circuits are kept to initialize the next iteration, with h being a hyperparameter. This procedure can be readily adapted to restrictions in hardware connectivity,

Algorithm 1 Cost function for classification accuracy of product-state training data over tensor elements.

Require: Input tensors $G_m, D_m \in \text{St}(\chi^d, \chi)$, $C \in \text{St}(\chi, N_C)$, $R \in S^X$.

Require: Training set of product states $\mathcal{T} = \{|\mathbf{x}_m\rangle, \ell_m, m = 1, \dots, M\}$

Require: Hyperparameters $N_b \geq 1$, $\chi \geq 1$.

function CLASSIFICATIONCOST($R, G, D, C, N_b, \chi, \mathcal{T}$)

$G \leftarrow \text{Reshape}(G_m, (\chi, d, \chi))$

$D \leftarrow \text{Reshape}(D_m^{\dagger}, (\chi, d, \chi))$

$B_{\alpha}^i \leftarrow \sum_{\beta} G_{\alpha\beta}^i R_{\beta}$

$V_{\alpha, \alpha'} \leftarrow \sum_i B_{\alpha}^{i*} B_{\alpha'}^i$

$N \leftarrow 1$

while $N < N_b'$ **do**

$B_{\alpha\beta}^i \leftarrow \sum_{\alpha'} V_{\alpha\alpha'} G_{\alpha'\beta}^i$

$V_{\alpha\alpha'} \leftarrow \sum_{\beta, i} B_{\beta\alpha'}^i G_{\beta\alpha}^{i*}$

$N \leftarrow N + 1$

end while

$B_{\alpha\beta}^m \leftarrow \sum_i x_{m,L-1}^i D_{\alpha\beta}^i$

$Q_{\alpha\beta}^m \leftarrow \sum_{\gamma} B_{\alpha\gamma}^m V_{\gamma\beta}$

$V_{\alpha\beta}^m \leftarrow \sum_{\gamma} Q_{\alpha,\gamma}^m B_{\beta\gamma}^{m*}$

$N \leftarrow 2$

while $N < L$ **do**

$B_{\alpha\beta}^m \leftarrow \sum_i x_{m,L-N}^i D_{\alpha\beta}^i$

$Q_{\alpha\beta}^m \leftarrow \sum_{\gamma} B_{\alpha\gamma}^m V_{\gamma\beta}^m$

$V_{\alpha\beta}^m \leftarrow \sum_{\gamma} Q_{\alpha,\gamma}^m B_{\beta\gamma}^{m*}$

$N \leftarrow N + 1$

end while

$\rho_{\ell\ell}^m \leftarrow \sum_{\alpha\beta} C_{\alpha\ell} V_{\alpha\beta}^m C_{\beta\ell}^*$

return $\frac{1}{M} \sum_m \mathcal{C}(\rho^m, \ell_m)$

end function

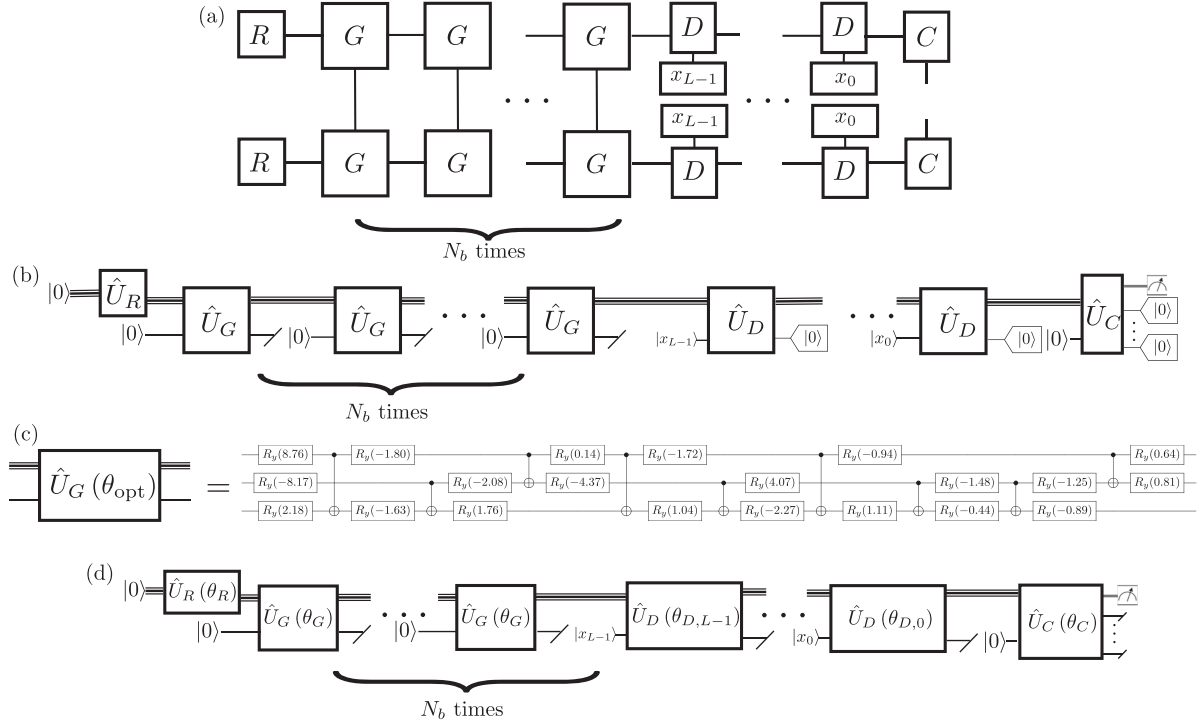


FIG. 3. Training of discriminator architecture. (a) Tensor-network diagram for the construction of $\rho_{\ell\ell}$ from the tensors $R, G, D,$ and $C,$ as in algorithm 1. (b) Compilation of the unitary embeddings of $R, G, D,$ and C results in a quantum circuit sampling from the distribution $\rho_{\ell,\ell}$ with postselection in the $|0\rangle$ state indicated by tags. (c) Example compiled unitary for \hat{U}_G evaluated at the optimal parameters for the postselected circuit in (b), here shown with two bond qubits. Gates connecting qubit wires are CNOTs, and the single-qubit rotations have the matrix representation $\hat{R}_y(\theta) \equiv \begin{pmatrix} \cos \frac{\theta}{2} & -\sin \frac{\theta}{2} \\ \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix}.$ (d) Circuit ansatz for model without postselection, in which the compiled circuit architectures are the same as the postselected model [e.g., panel (c) for \hat{U}_G], but with different values for the parameters.

can be modified to account for hardware noise through the cost function used for model downselection, and utilizes the freedom in compiling the isometric tensors to unitaries to discover compiled gate sequences that are as short as possible.

As also described in Ref. [34], even when the MPS tensors are in a properly defined canonical form there is still a unitary gauge freedom on the bond degrees of freedom that can be exploited to reduce the complexity of the compiled unitary embedding. In the finite-size systems studied in Ref. [34], a “diagonal gauge” was put forward in which a permutation operator was applied on the bond degrees of freedom to attempt to maximize diagonal dominance of the compiled unitary in the computational basis. In this work we define an analog of this diagonal gauge for iMPSs. Namely, we optimize a general unitary rotation on the bond degrees of freedom \mathbb{W} over the cost function

$$\sum_{\alpha,\beta} \Delta_{\alpha\beta} \left[\sum_i \tilde{G}_{\alpha\beta}^i \tilde{G}_{\alpha\beta}^{i*} + \tilde{D}_{\alpha\beta}^i \tilde{D}_{\alpha\beta}^{i*} \right], \tag{15}$$

in which $\Delta_{\alpha,\beta} = \sum_{\mu} |b_{\mu}(\alpha) - b_{\mu}(\beta)|^2$ is a matrix of differences between the binary vector representations $\mathbf{b}(\gamma)$ (i.e., Hamming distance) of the bond indices that penalizes off-diagonal terms and

$$\tilde{G}_{\alpha\beta}^i = \sum_{\alpha'\beta'} W_{\alpha\alpha'} G_{\alpha'\beta'}^i W_{\beta\beta'}^*, \tag{16}$$

$$\tilde{D}_{\alpha\beta}^i = \sum_{\alpha'\beta'} W_{\alpha\alpha'} D_{\alpha'\beta'}^i W_{\beta\beta'}^*. \tag{17}$$

As above, we optimize this cost function using Riemannian optimization over the elements $\mathbb{W} \in \text{St}(\chi, \chi).$ Following this optimization, the tensors $R, G, D,$ and C have their bond degrees of freedom transformed using this unitary operator. An example of a gauged and compiled circuit using our compilation procedure is shown in Fig. 3(c) using CNOTs and R_y rotations.

The compiled unitaries embedding the tensors $R, G, D,$ and C immediately define a quantum circuit which produces samples of the binary encoding of the class label ℓ according to the distribution $\rho_{\ell,\ell}.$ However, because of the isometric nature of the tensors R and others, this circuit requires postselection of the physical qubit from \hat{U}_D and the bond qubits from $\hat{U}_C,$ as shown in Fig. 3(b). We can remedy this by replacing all instances of postselection by a trace over the corresponding qubits; however, the parameters of the model have been optimized for the postselected operation by virtue of the isometric nature of, e.g., the tensors $R,$ and using these same parameters directly without postselection usually results in poor performance. One way to improve performance is to treat each of the operators, e.g., $\hat{U}_G,$ as a general unitary and optimize its parameters for the circuit without postselection. Such an approach has been taken in Ref. [22],

where the unitaries were parameterized as the exponential of a Hermitian matrix and the elements of this Hermitian matrix optimized by the simultaneous perturbation stochastic approximation (SPSA). However, such an approach will produce dense unitaries which then have to be compiled to quantum hardware, generally requiring deep circuits [52]. In contrast, our approach produces underconstrained isometries, and the freedom in embedding these isometries into unitaries can be used to produce shorter circuits accounting for hardware constraints.

Our method to improve the performance of the model architecture when applied without postselection is to keep the circuit architecture that has been learned for operation with postselection [e.g., the circuit compiling the unitary embedding \hat{U}_G in Fig. 3(c)], and to directly optimize the parameters of this circuit for a cost function using operation without postselection. This leads to the circuit architecture in Fig. 3(d) with angles θ as the parameters to be optimized. As the operation of this circuit is now a completely positive trace-preserving (CPTP) map, $\sum_\ell \rho_{\ell\ell} = 1$. In order to balance the optimization across many training instances in the dataset, we find it is useful to use the regularized cost function of Ref. [22]

$$C_{\lambda,\eta}(\rho, \ell_m) = \max(\rho_{\bar{\ell}_m \bar{\ell}_m} - \rho_{\ell_m \ell_m} + \lambda, 0)^\eta, \quad (18)$$

in which λ and η are hyperparameters and $\bar{\ell}_m$ is the label of the incorrect class for training vector $|\mathbf{x}_m\rangle$ with highest model probability. This results in algorithm 2 for the classification accuracy cost given a set of circuit parameters $\theta = \{\theta_R, \theta_G, \{\theta_{D,i}, i = 0, \dots, L-1\}, \theta_C\}$.

In practice, we optimize the circuit using the Limited Memory Broyden-Fletcher-Goldfarb-Shanno (LBFGS) algorithm implemented in PyTorch [53].

B. Model inference with entangled data

Once the optimal parameters θ_{opt} for execution of the model without postselection have been determined, Algorithm 2 can be simply adapted to inference the model by replacing the final cost function evaluation with sampling ℓ from $\rho_{\ell\ell}$. As described, this algorithm applies to unentangled product state data, but this same procedure can be readily adapted to inference of entangled data. Here we present an algorithm for inferencing entangled data in the form of an iMPS model. Namely, the test data are assumed to be provided in the form of unitaries \hat{U}_R and \hat{U}_G embedding a right boundary condition and a left-canonical MPS tensor A , respectively. An algorithm to obtain this form for the ground state of a Hamiltonian with known matrix product operator (MPO) form is presented in Appendix A.

As described in Sec. II, the procedure to generate a finite state string from a left canonical iMPS is to first obtain the half-infinite density matrix by iteration of the transfer operator on the right boundary vector, and then to use the iMPS tensor with this density matrix as input to generate samples. In practice, we will obtain the half-infinite density matrix using a finite number N'_b of iterations as in Eq. (5), with N'_b a hyperparameter (set to N_b in the results that follow). Hence, the process for inferencing entangled state is (1) generate the half-infinite density matrix of the test state via N'_b iterations of \hat{U}_G applied to the right boundary bond state generated by \hat{U}_R

Algorithm 2 Cost function for classification accuracy of product-state training data without postselection.

Require: Input parameterized gate sequences $\hat{U}_R(\bullet)$, $\hat{U}_G(\bullet)$, $\hat{U}_D(\bullet)$, $\hat{U}_C(\bullet)$ and initial guesses for the parameters $\theta = \{\theta_R, \theta_G, \{\theta_{D,i}, i = 0, \dots, L-1\}, \theta_C\}$.

Require: Training set of product states $\mathcal{T} = \{|\mathbf{x}_m\rangle, \ell_m, m = 1, \dots, M\}$

function CLASSIFICATIONCOSTNOPOSTSELECTION($U_R, U_G, U_D, U_C, N_b, \theta, \mathcal{T}$)

$B_{\alpha}^i \leftarrow \sum_{\beta} U_{G;(\alpha,i)(\beta,0)}(\theta_G) U_{R;\beta,0}(\theta_R)$

$V_{\alpha,\alpha'} \leftarrow \sum_i B_{\alpha}^{i*} B_{\alpha'}^i$

$N \leftarrow 1$

while $N < N'_b$ **do**

$B_{\alpha\beta}^i \leftarrow \sum_{\alpha'} V_{\alpha\alpha'} U_{G;(\alpha',i)(\beta,0)}(\theta_G)$

$V_{\alpha\alpha'} \leftarrow \sum_{\beta,i} B_{\beta\alpha'}^i U_{G;(\alpha,i)(\beta,0)}^*(\theta_G)$

$N \leftarrow N + 1$

end while

$B_{\alpha\beta}^{mp} \leftarrow \sum_i x_{m,L-1}^i U_{D;(\alpha,p)(\beta,i)}(\theta_{D,L-1})$

$Q_{\alpha\beta}^{mp} \leftarrow \sum_{\gamma} B_{\alpha\gamma}^{mp} V_{\gamma\beta}$

$V_{\alpha\beta}^m \leftarrow \sum_{\gamma,p} Q_{\alpha,\gamma}^{mp} B_{\beta\gamma}^{mp*}$

$N \leftarrow 2$

while $N < L$ **do**

$B_{\alpha\beta}^{mp} \leftarrow \sum_i x_{m,L-N}^i U_{D;(\alpha,p)(\beta,i)}(\theta_{D,L-N})$

$Q_{\alpha\beta}^{mp} \leftarrow \sum_{\gamma} B_{\alpha\gamma}^{mp} V_{\gamma\beta}^m$

$V_{\alpha\beta}^m \leftarrow \sum_{\gamma,p} Q_{\alpha,\gamma}^{mp} B_{\beta\gamma}^{mp*}$

$N \leftarrow N + 1$

end while

$\rho_{\ell\ell}^m \leftarrow \sum_{\alpha,\beta,\gamma} U_{C;(\gamma,\ell),(\alpha,0)}(\theta_C) V_{\alpha\beta}^m U_{C;(\gamma,\ell),(\beta,0)}^*(\theta_C)$

return $\frac{1}{M} \sum_m C_{\lambda\eta}(\rho^m, \ell_m)$

end function

on a register of $\log_2 \chi'$ bond qubits $|b'\rangle$; (2) generate the half-infinite density matrix of the feature space employed by the classifier by applying N_b iterations of $\hat{U}_G(\theta_G)$ applied to the right boundary bond state generated by $\hat{U}_R(\theta_R)$ on a different register of $\log_2 \chi$ bond qubits $|b\rangle$; (3) for $i = 1, \dots, L$, generate a qubit of the test state using \hat{U}_G applied to the register $|b'\rangle$ and a single physical qubit in the $|0\rangle$ state, feed this state into the classifier by applying $\hat{U}_D(\theta_{D:L-i})$ to the prepared state and the classifier bond register $|b\rangle$ and resetting the physical qubit; (4) apply the classifier unitary $\hat{U}_C(\theta_C)$ to the classifier bond register $|b\rangle$; and (5) measure the subset of qubits defining the class register and discard all others. This process is shown graphically in Fig. 4. The qubit resource requirements, provided the same physical qubit is utilized by both the test state and the discriminator to generate their half-infinite density matrices, are $\log_2 \chi + \log_2 \chi' + 1$. Utilizing two physical qubits—one for the test state and one for the discriminator—leads to the resource requirement $\log_2 \chi + \log_2 \chi' + 2$, but enables the parallelized generation of the half-infinite density matrices shown in Fig. 4.

IV. DEMONSTRATION ON THE TRANSVERSE FIELD ISING MODEL

In this section we describe an example application of the discriminator architecture, training, and test strategies put forth in Sec. III. In particular, we consider the problem of discriminating the two phases of the spin-1/2 transverse field Ising model (TFIM) in one dimension, a canonical model

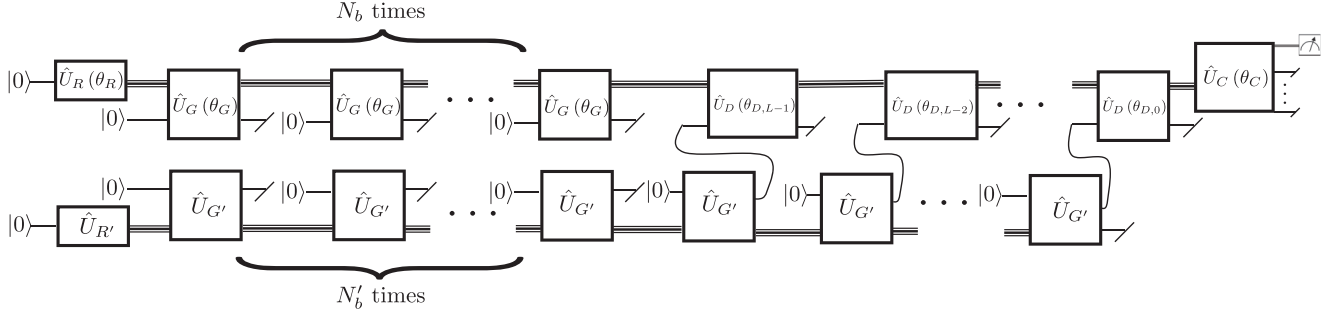


FIG. 4. Inference of an MPS with the discriminator architecture. The diagram shows inference of an iMPS test state (lower half of diagram) with unitaries $\hat{U}_{G'}$ and $\hat{U}_{R'}$ using the discrimination circuit with unitaries \hat{U}_R , \hat{U}_G , \hat{U}_D , and \hat{U}_C (top half of diagram). The test state generates its half-infinite density matrix using N_b' iterations of the transfer operator encoded by $\hat{U}_{G'}$ applied to the boundary state generated by $\hat{U}_{R'}$, while the discriminator generates the feature space half infinite density matrix using N_b iterations of \hat{U}_G applied to the boundary state generated by \hat{U}_R on a different bond register. The test state generator then generates a single qubit of the test state at a time, this state is fed into the discriminator to condition its bond register, and the physical qubit is reset. After L qubits have been generated of the state, the test state bond qubits are discarded, the discriminator bond qubits are fed into the classifier circuit \hat{U}_C , and the measurement of the class register qubits defines the predicted class label ℓ .

for quantum phase transitions [54] that has been realized in a variety of physical systems [55–57]. We demonstrate this application in both classical simulations and on Quantinuum’s H1-2 trapped ion quantum computer. The Hamiltonian for the antiferromagnetic TFIM is

$$\hat{H}_{\text{TFIM}} = \sum_{\langle i,j \rangle} \hat{\sigma}_i^z \hat{\sigma}_j^z - h \sum_i \hat{\sigma}_i^x, \quad (19)$$

in which $\langle i, j \rangle$ denotes a sum over nearest-neighbor pairs i, j and we have normalized units of energy to the antiferromagnetic coupling strength. For $h \gg 1$ the ground state is in a paramagnetic (PM) phase whose characteristics are captured by the classical state $|\rightarrow \cdots \rightarrow\rangle$ with all spins pointing along the x direction. When $h \ll 1$ the ground state is in the antiferromagnetic (AFM) phase with properties similar to the classical states $|\uparrow\downarrow \cdots \uparrow\downarrow\rangle$ and $|\downarrow\uparrow \cdots \downarrow\uparrow\rangle$. There is a quantum phase transition between the PM and AFM phases at $h = 1$ in the thermodynamic limit, and we will train our discrimination architecture on data taken on either side of the transition with the goal of discriminating the phase of an unknown test quantum state.

A. Description of dataset

Our dataset will comprise single-shot measurements taken from the ground state of Eq. (19) at some value of h . We will define a single-shot measurement through a collection of L Hermitian operators $\{\hat{O}_i\}$, $i = 1, \dots, L$ which act on a subset of L contiguous sites \mathcal{L} from the infinite lattice indexed by i . Each of these operators admits a spectral decomposition as $\hat{O}_i = \sum_{\lambda_i} \lambda_i |\lambda_i\rangle\langle\lambda_i|$. A single-shot measurement is now defined as a simultaneous projective measurement of all of the operators $\{\hat{O}_i\}$, which we can denote as a collection of the associated eigenvalues $\{\lambda_i\}$. Associated with this measurement is a classical product state (CPS) $|\lambda_1, \dots, \lambda_L\rangle$, which is the state produced by the projective measurement. The association is provided by the fact that the measurement $\{\lambda_i\}$ from a state $|\psi\rangle$ occurs with probability $\text{Tr}[|\lambda_1, \dots, \lambda_L\rangle\langle\lambda_1, \dots, \lambda_L| \text{Tr}_{i \notin \mathcal{L}} |\psi\rangle\langle\psi|]$. A dataset

comprises a collection of these CPSs, randomly chosen from the ground state $|\psi\rangle$ at specific points within the phase diagram, together with labels ℓ denoting the quantum phase of the state at that point, $\mathcal{T} = \{|\lambda_m\rangle, \ell_m, m = 1, \dots, M\}$.

We generate our single-shot measurement training data using MPS methods. Namely, we find the ground state at transverse field strength h on $F \gg L$ sites using standard variational MPS methods [11], and then sample CPSs from this MPS given an operator basis $\{\hat{O}_i, i \in \mathcal{L}\}$ on a subset of L sites \mathcal{L} . In the results presented in this work, we take $F = 32$, a truncated weight of $\varepsilon = 10^{-6}$, maximum bond dimension of $\chi_{\text{max}} = 40$, $L = 6$, and \mathcal{L} to be the L sites in the center of the chain. We take half of our samples in the uniform z basis $\{\hat{O}_i = \hat{\sigma}_i^z, \forall i \in \mathcal{L}\}$ and half in the uniform x basis $\{\hat{O}_i = \hat{\sigma}_i^x, \forall i \in \mathcal{L}\}$. These choices reflect how we envision our discrimination architecture could be used in a near-term setting; the central region of a finite-size experimental system will be measured and taken to be representative of the bulk of an infinite system. Present-day quantum simulators, e.g., those based on trapped cold atoms, do not universally feature single-site addressability, but many have the ability to change the global measurement basis, e.g., by applying a uniform microwave pulse. Generating uncorrelated CPS samples from an MPS has been described before in the context of minimally entangled typical thermal state (METTS) algorithms and is reviewed in Appendix B.

We take the states defining the PM and AFM phases to have $h = 10$ and $h = 0.1$, respectively, and generate 1000 shots at each value of h in both the x and z basis. These shots are randomly reshuffled using an 80/20 split into a training set and a test set. We note that the ground state in the AFM phase is doubly degenerate, with the ground states being distinguished by the global spin flip operator $\hat{X} = \prod_{i=1}^F \hat{\sigma}_i^x$ that generates a \mathbb{Z}_2 symmetry. Numerical noise in the initialization of the MPS ground state search procedure can result in finding a symmetry-broken state in this manifold. If all training measurements were taken from a single MPS representation, this would result in significant bias compared to measurements taken from an experiment, in which a different symmetry

broken state could be obtained for each realization. In order to avoid this phenomenon, we add a term $-h_{\mathbb{Z}_2}\hat{X}$ to the Hamiltonian with $h_{\mathbb{Z}_2} \sim 10$ to ensure that the \mathbb{Z}_2 -symmetric ground state is chosen and there is no symmetry-breaking bias in the training data.

Finally, for entangled inferencing of a model trained on the product state data just described (see Sec. III B), we also generate iMPS models of the ground state with small bond dimension $\chi \ll \chi_{\max}$ through a novel optimization procedure described in Appendix A. This algorithm produces a translationally invariant iMPS tensor A' in left-canonical form together with a boundary bond state R' that optimizes the production of the true half-infinite density matrix obtained from $|\psi[A']\rangle$ from a finite “burn-in” procedure with a fixed number N'_b of iterations. Following optimization, these tensors can be gauged and compiled as described in Sec. III A to define circuits for inferencing as described in Sec. III B.

B. Training and evaluation

Here we apply the classical training and compilation procedures of Sec. III A to the dataset described in Sec. IV A. Our main metric of performance will be the F_1 score, defined from precision p and recall r as

$$p_i = \frac{C_{ii}}{\sum_j C_{ji}}, \quad (20)$$

$$r_i = \frac{C_{ii}}{\sum_j C_{ij}}, \quad (21)$$

$$[F_1]_i = 2 \frac{p_i r_i}{p_i + r_i}. \quad (22)$$

In these expressions, \mathbb{C} is the confusion matrix indexed by classes whose element C_{ij} represents the number of data elements predicted to be in class j whose truth class is i .

As a classical machine learning benchmark, we train a random forest classifier using the implementation in scikit-learn [58] with the hyperparameter `n_estimators` set to 20. Here the classical product state single-shot measurements are flattened into classical vectors of dimension $2L$ as

$$|\mathbf{x}\rangle = \left(\sum_{i_0} x_{0,i_0} |i_0\rangle \right) \otimes \cdots \otimes \left(\sum_{i_{L-1}} x_{L-1,i_{L-1}} |i_{L-1}\rangle \right) \quad (23)$$

$$\rightarrow (x_{0,0}, x_{0,1}, \dots, x_{L-1,0}, x_{L-1,1})^T, \quad (24)$$

and these classical representations are used for learning. This flattened representation can capture superpositions at the single qubit level, but not entanglement between qubits. We note that no claims are made about the optimality of this classical data encoding or classifier, we are merely using it as a reasonable classical benchmark.

Following training of the tensor elements using manifold gradient descent on the cost function in algorithm 1, we find class-averaged F_1 scores of $\bar{F}_1 = 0.814, 0.910,$ and 0.939 on the training data for models with $\chi = 2, 4,$ and 8 , corresponding to one, two, and three bond qubits, respectively. The classical random forest benchmark achieves a class-averaged F_1 score of 0.973 on the training set. We then compile the associated unitary embeddings using the greedy heuristics described in Sec. III A with a squared two-norm tolerance of 4×10^{-4} , a basis of CNOT entangling gates and R_y single-qubit

TABLE I. CNOT counts for the compiled unitaries with increasing bond dimension.

	$\chi = 2$	$\chi = 4$	$\chi = 8$
\hat{U}_R	0	1	3
\hat{U}_G	2	9	42
\hat{U}_D	2	9	41
\hat{U}_C	2	3	8

rotations, and without any restrictions on the qubit pairs CNOTs can be applied to. The CNOT counts for the various compiled unitaries are collected in Table I.

While the finite tolerance of compiling the unitaries does not significantly impact the model when run with postselection, the \bar{F}_1 scores using these parameters without postselection are $0.360, 0.364,$ and 0.228 . Following 30 epochs of LBFGS optimization of the compiled unitary parameters with the cost function hyperparameters $\lambda = 0.9, \eta = 2.0$, we find optimized \bar{F}_1 scores of $0.764, 0.822,$ and 0.860 . In addition to the \bar{F}_1 scores, an important metric as χ increases is the single-shot probability of success $P_{SS} = \rho_{\ell_m, \ell_m}$, which determines the frequency with which the classifier will return the correct label in a single execution of the circuit. As shown in Fig. 5, increasing the bond dimension not only increases the average number of correct classifications (captured by the \bar{F}_1 score), but also the probability that a single run of the circuit produces the correct label, on average. We note that the clustering of highest probabilities near 0.9 is influenced by the choice of the hyperparameter λ ; different choices of hyperparameters, possibly bond-dimension-dependent, can alter this distribution.

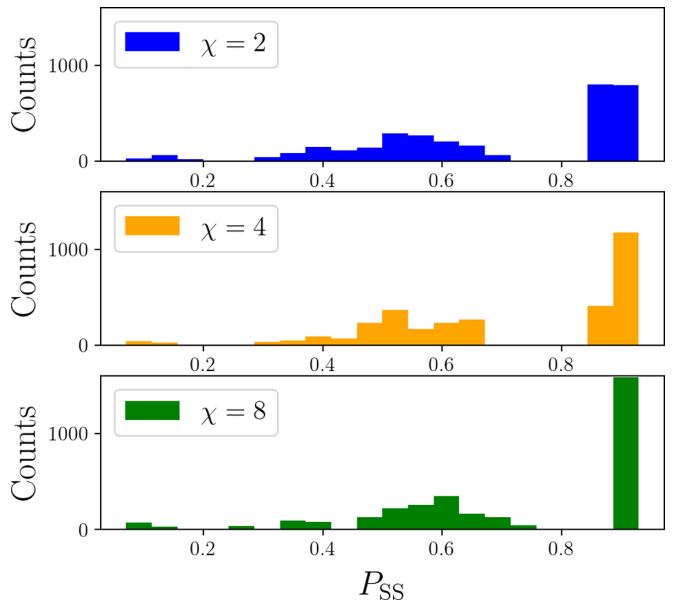


FIG. 5. Dependence of single-shot probability of success on bond dimension. Histograms of the single-shot probability of correct classification P_{SS} over the training set are shown for $\chi = 2, 4, 8$. As the bond dimension increases, the average probability of correct classification also increases.

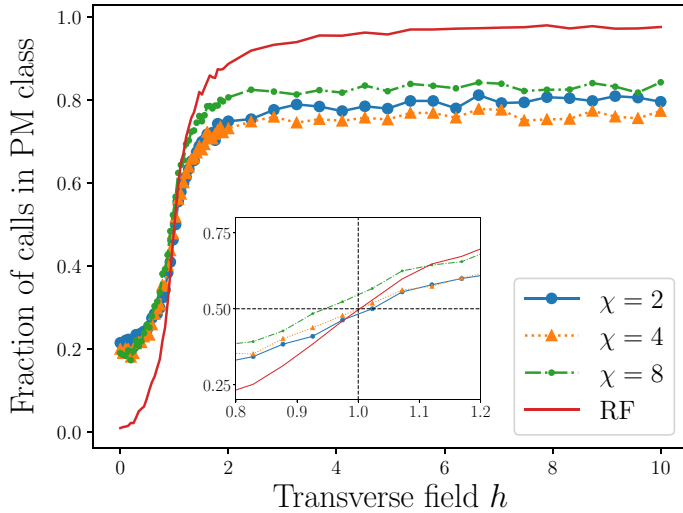


FIG. 6. Predicted phases from the tensor-network discriminator and classical benchmark; product state input data. The probability that a classifier predicts the PM class, averaged over the test data, is shown as a function of the transverse field h . The inset shows a zoom around the true phase transition point, $h = 1$, demonstrating that all classifiers have a bias of a few percent.

A canonical use case for a quantum classifier is determining the quantum phase of the state producing the measurements, where the classifier has been trained on data deep in the phases to be distinguished. In this use case, a key metric is the distance between the phase transition point predicted by the classifier and the true transition point. In the case of binary classification, the predicted phase transition point can be taken as the place where the classifier predicts 50% probability for both phases. In order to test this use case, we produced ground state representations in the range $h \in [0, 10]$ using MPS variational ground state search and sampled 1000 single-shot measurements in the x and z basis for each h , following the procedure described in Sec. IV A. The fraction of data predicted to be in the PM class for the $\chi = 2, 4, 8$ tensor-network discriminators and the random forest benchmark are shown in Fig. 6. We see slight improvements in the probability of correct classification as the bond dimension is increased, in accordance with the behavior seen in the training set in Fig. 5, while the random forest classifier shows the best performance by this metric. As seen in the inset, all classifiers are remarkably unbiased in their determination of the phase transition point at $h = 1$ even though they were trained on data at $h = 0.1$ and $h = 10$; all classifiers are within a few percent of this value.

The final use case we consider is inferencing of our models on entangled input data. To generate the test data for this use case, we produce iMPS models for the ground state at a variety of transverse field strengths h as described in Appendix A. We stress that this process is different than the finite-size ground state search used to generate the single-shot measurement training and test data. We then inference our tensor-network models using a string of $L = 6$ qubits of entangled data generated from these iMPS models as described in Sec. III B. The results are shown in Fig. 7. Here we plot the probability that the classifier determines the input state to be in the PM




FIG. 7. Predicted phases from the tensor-network discriminator and classical benchmark; entangled input data. The probability of a classifier returning a call in the PM class for a single run of the circuit is shown as a function of the transverse field h . As the bond dimension of the classifier increases, the probability of correct classification increases. The performance of all classifiers is relatively unchanged as the bond dimension of the iMPS generating the test data is increased from $\chi_i = 4$ to $\chi_i = 8$.

class in a single run of the circuit, as opposed to the case of Fig. 6, in which probabilities were inferred by averaging the 2000 elements of the test set. Increasing the bond dimension of the classifiers improves their performance in the sense that the single-run probability of correct classification increases. We see little dependence of the performance on the bond dimension of the input data, χ_i . This process of entangled inference following training of the classifier on single-shot measurements provides a metric for determining the consistency of experimental data (single-shot measurements) with a reduced-order model (entangled input data according to a particular model Hamiltonian). Finally, we note that there is no classical benchmark here, as the inferencing of entangled data is an inherently quantum process. It is interesting to note that the performance of the classifiers on entangled input data even at the smallest bond dimension is higher than the classifiers evaluated on single-shot data at any of the evaluated bond dimensions, even though the classifiers were trained only on single-shot measurements.

C. Hardware implementation

We now turn to results obtained on Quantinuum's H1-2 trapped ion quantum computer [59]. A key consideration when running the model on quantum hardware is the number of runs of the circuit required to obtain a high-confidence result for the predicted class. Each run of the circuit with data instance $|\mathbf{x}\rangle$ returns a class label ℓ which is 0 with an unknown probability $p_0(\mathbf{x})$. The task of the classical postprocessing is to infer what $p_0(\mathbf{x})$ is from a finite number n of runs of the circuit, and assign the class of \mathbf{x} as 0 if $p_0(\mathbf{x}) > 0.5$ and 1 if $p_0(\mathbf{x}) < 0.5$. In the ideal case where hardware noise is independent and identically distributed, each run of the circuit constitutes a Bernoulli trial, and we can utilize

binomial confidence intervals to estimate the uncertainty in the model’s predictions. In what follows, we will employ the Wilson score interval, which utilizes a normal approximation to the binomial to derive closed-form expressions for the confidence interval in terms of z , the $1 - \frac{\alpha}{2}$ quantile of the normal distribution with target error rate α . Namely, the Wilson score interval predicts

$$p_0(\mathbf{x}) \approx \frac{n_0 + z^2/2}{n + z^2} \pm \frac{z}{n + z^2} \sqrt{\frac{n_0 n_1}{n} + z^2/4}, \quad (25)$$

where we have suppressed the dependence of n_0 , n_1 , and n on \mathbf{x} on the right-hand side to keep the notation uncluttered. In the results presented, we use the 90% confidence interval, corresponding to $z = 1.645$. A complementary online algorithm which determines when a sufficient number of shots has been taken according to a Bayesian credible interval is given in Appendix C.

Given the better performance of the models inferred on entangled input data vs single-shot data in classical simulation (see Sec. IV B), we utilized the entangled inference strategy on the quantum hardware. Also noting the relative insensitivity of our model performance to the bond dimension of the entangled input data (see Fig. 7), we utilized a fixed input data bond dimension of $\chi_i = 4$ for all experiments. For each value of the transverse field h considered, the circuit was run for 100 shots, and the Wilson score interval utilized to obtain the uncertainty due to a finite number of shots. In order to identify the predicted transition point, we take the points closest to 50% predicted probability in the PM phase and fit them to a linear regression model with the Wilson score intervals as errors. Figures 8 and 9 show the results for $\chi = 2$ and 4, respectively. In both figures, the orange triangles correspond to the fraction of shots classifying the input state as being in the PM class, the blue circles are the same quantity computed using an H1 simulator in the absence of hardware noise, the error bars are the 90% Wilson confidence intervals, and the black dashed line is the prediction of the fit linear regression model. Identifying the point of “maximal confusion” of the classifier as the transition point, the linear regression predicts $h = 0.962$ and 0.994 for the $\chi = 2$ and $\chi = 4$ models, respectively. The mean absolute error (MAE) and R^2 coefficient of determination of the two linear models over the data used in the regression are (0.03,0.90) and (0.04,0.80) for $\chi = 2$ and 4, respectively, indicating that the lower bond dimension data is slightly better fit and the linear model captures more of the variance in those data. This is to be expected on account of the deeper circuits utilized for the $\chi = 4$ model.

V. CONCLUSIONS AND OUTLOOK

We presented a workflow for classifying a translationally invariant quantum state based on L qubits of quantum data extracted from it using a tensor-network (TN) architecture with a one-dimensional matrix product state (MPS) geometry. In our discriminator architecture, a register of $\log_2 \chi$ ancilla qubits forms an effective bond space of dimension χ mediating correlations between physical degrees of freedom. Similar to encoder methods in classical machine learning, this bond space stores an effective reduced-dimensional latent representation of the data the model represents. Inferencing of

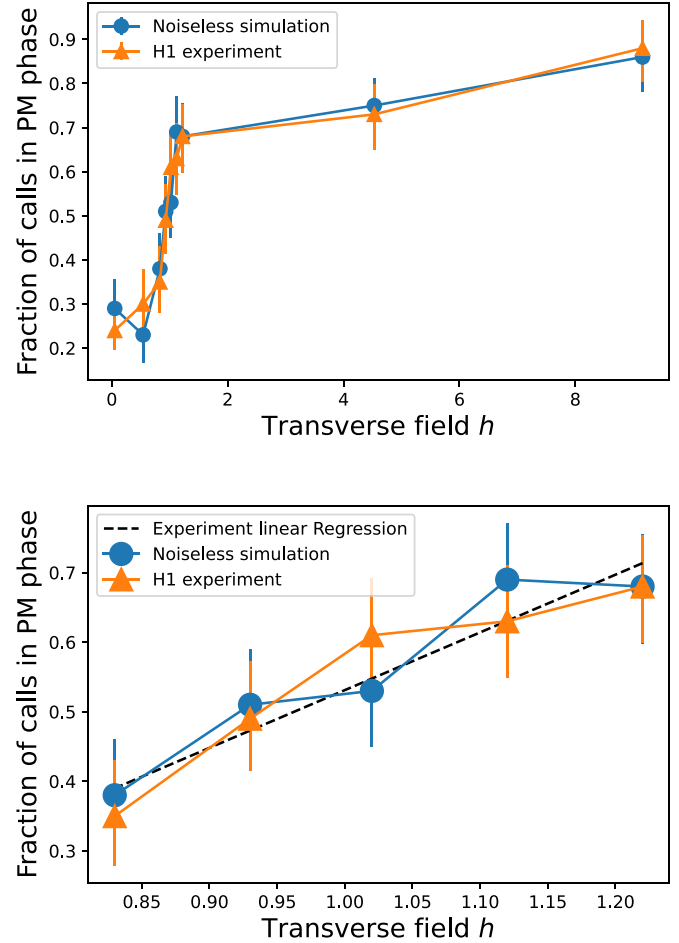


FIG. 8. Performance of $\chi = 2$ model on entangled data. The top panel shows the predictions of the $\chi = 2$ tensor-network discriminator evaluated on entangled iMPS input data with $\chi_i = 4$, with the triangles corresponding to the mean of experimental results on the H1 machine and the circles corresponding to simulations of the circuit execution in the absence of hardware noise. All points correspond to 100 shots, and error bars are 90% Wilson confidence intervals. The bottom panel is a zoom around the transition point, together with the linear regression model predicting a transition at $h = 0.962$.

our model begins with preparing the ancilla qubits in an initial state which encodes a prior distribution over a collection of latent quantum feature vectors. A series of quantum operations coupling these ancilla qubits with a register of physical qubits then conditions the state of the ancillas on the input quantum data. A final quantum operation predicts the class of the input data depending on the state of the ancilla qubits and stores the result in a subset of the ancillas. We presented methods for robust and efficient classical preconditioning of this discriminator architecture from a dataset of single-shot experimental measurements, consisting of optimization of the isometric tensors in a classical tensor-network representation using manifold optimization techniques, compilation of the isometric tensors into quantum operations using hardware-aware greedy compilation heuristics, and further optimization of the parameters of the resulting compiled quantum circuit to remove the postselection requirements of the isometric tensor-network model. Following training, the resulting model

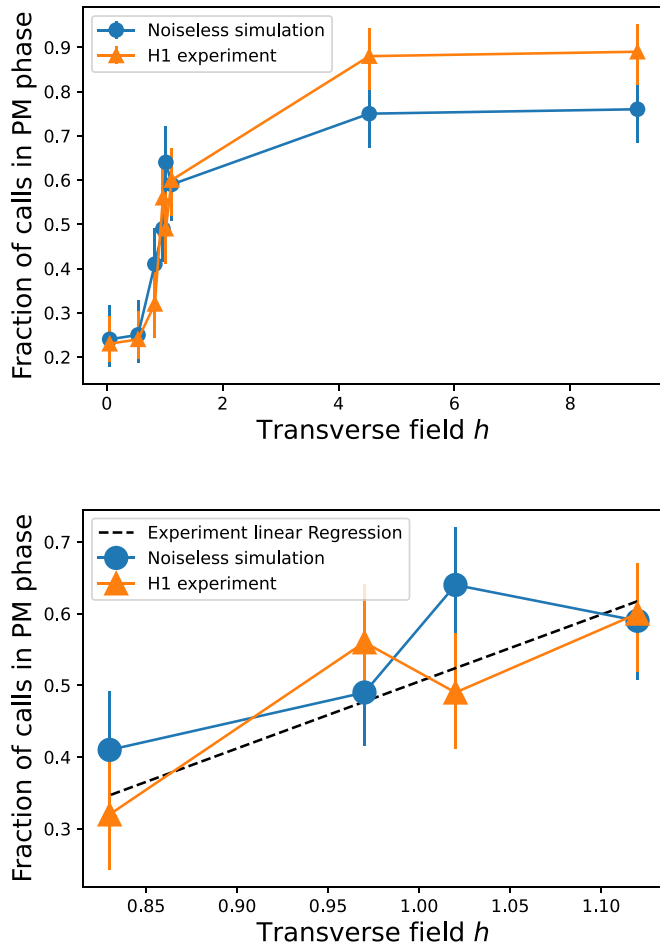


FIG. 9. Performance of $\chi = 4$ model on entangled data. The top panel shows the predictions of the $\chi = 4$ tensor-network discriminator evaluated on entangled iMPS input data with $\chi_i = 4$, with the triangles corresponding to the mean of experimental results on the H1 machine and the circles corresponding to simulations of the circuit execution in the absence of hardware noise. All points correspond to 100 shots, and error bars are 90% Wilson confidence intervals. The bottom panel is a zoom around the transition point, together with the linear regression model predicting a transition at $h = 0.994$.

can be inferred with quantum data on quantum hardware. Our architecture enables the discovery of model circuits in a data-driven fashion without specifying the gate depth or composition ahead of time and aids in trainability by enabling efficient classical preconditioning and requiring measurements of only a small subset of model qubits.

We demonstrated the training of our TN discriminator on a synthetic dataset of data extracted from the 1D transverse field Ising model (TFIM) in its paramagnetic (PM) and antiferromagnetic (AFM) phases. Following training on single-shot experimental data, we analyzed the performance of the model inferred on product state or entangled input data. We found that increasing the bond dimension of the model, amounting to adding more qubits to the ancilla register in the quantum representation, improved both the average classification accuracy and the probability of correct classification in a single run of the circuit, and that inference of entangled data had better average performance than inference of unentangled data. Using

TN models of bond dimension $\chi = 2$ and 4 trained on single-shot data deep in the PM and AFM phases, we ran experiments on Quantinuum’s model H1-2 trapped ion quantum computer inferring these models against entangled input data modeled as MPSs of bond dimension $\chi_i = 4$ across the phase diagram. Using linear regression near the point of maximal confusion of our models (i.e., 50% prediction in each phase), we extracted predictions of the phase transition at $h = 0.962$ and 0.994 for bond dimension $\chi = 2$ and $\chi = 4$, comparing well with the known location of $h = 1$.

Our results demonstrate that classical techniques for optimizing TNs can be utilized to define and optimize well-performing machine learning models for quantum hardware in a data-driven and hardware-aware fashion. The classical training and optimization schemes put forwards in this paper can be used to precondition models with a modest number of qubits to reduce the workload on an optimizer when operated with a quantum device in the loop. Because tensor-network model expressibility, as quantified by the bond dimension, grows exponentially with the number of qubits but only polynomially with classical resources, adding additional qubits provides a clear pathway for potential model advantage. By analyzing a series of classically preconditioned model architectures of increasing bond dimension, model architectures can be scaled towards the classically intractable regime in which quantum models may display an advantage in expressibility. While we focused on the case of MPS models, the generalization to other TN topologies is straightforward. In addition, looking beyond machine learning, many of the tools we have developed can be brought to bear on efficiently preconditioning and scaling quantum TN models for other use cases, such as quantum simulation and error correction [60,61].

ACKNOWLEDGMENTS

We would like to thank Matt Abernathy and Giuseppe D’Aguanno for collaboration on related work. M.L.W., P.T., and G.Q. would like to acknowledge funding from the Internal Research and Development program of the Johns Hopkins University Applied Physics Laboratory. K.R.A.H acknowledges support from the Welch Foundation through Grant No. C1872, the Office of Naval Research through Grant No. N00014-20-1-2695, and the National Science Foundation through Grants No. PHY1848304 and No. CMMI2037545. K.R.A.H. also benefited from discussions at the KITP, which was supported in part by the National Science Foundation under Grant No. NSF PHY-1748958.

APPENDIX A: ALGORITHM FOR OPTIMIZING AN IMPS IN LEFT-CANONICAL FORM

In this Appendix we provide an algorithm for obtaining a quantum circuit for the iMPS description of the ground state of a Hamiltonian \hat{H} with a known matrix product operator (MPO) form. While the iMPS algorithm put forth by McCulloch [62] (building on the seminal iDMRG algorithm of White [20]) works well for this purpose, the algorithms herein provide a simple and robust means for obtaining iMPS representations directly in left-canonical form, from which

our quantum compilation methods can be applied to generate low-depth quantum circuits.

We define the action of the generalized left transfer operator on a $\chi \times \chi$ matrix \mathbb{E} as

$$T_{L,\hat{O}}(A; \mathbb{E}) = \sum_{i,i'} \langle i | \hat{O} | i' \rangle \mathbb{A}^{i\dagger} \mathbb{E} \mathbb{A}^{i'}, \quad (\text{A1})$$

where \hat{O} is an operator acting on the single-site Hilbert space and A is the translationally invariant iMPS tensor. Hence, the left transfer operator of the identity for an iMPS in left-canonical form has the $\chi \times \chi$ identity matrix \mathbb{I} as an eigenmatrix with eigenvalue 1, $T_{L,I}(A; \mathbb{I}) = \mathbb{I}$. Additionally, defining the corresponding right generalized transfer matrix

$$T_{R,\hat{O}}(A; \mathbb{E}) = \sum_{i,i'} \langle i | \hat{O} | i' \rangle \mathbb{A}^i \mathbb{E} \mathbb{A}^{i\dagger}, \quad (\text{A2})$$

we see that the right transfer operator of the identity admits the half-infinite reduced density matrix as an eigenmatrix with eigenvalue 1 in the case of a left-canonical iMPS, $T_{R,I}(A; \rho_{\text{half-infinite}}) = \rho_{\text{half-infinite}}$. Defining a matrix representation of the transfer operator as

$$[T_{\hat{O}}(A)]_{(\alpha,\alpha'),(\beta,\beta')} = \sum_{ii'} A_{\alpha\beta}^{i*} \langle i | \hat{O} | i' \rangle A_{\alpha'\beta'}^{i'}, \quad (\text{A3})$$

we see that the left and right transfer operators are just the left- and right-action of this transfer operator, and so $T_I(A)$ admits \mathbb{I} and $\rho_{\text{half-infinite}}$ as left- and right-eigenmatrices with eigenvalue 1, respectively.

The transfer operator formalism is useful to extract fixed-point relations of the iMPS state, such as the eigenmatrix relations described in the last paragraph. For a Hamiltonian with MPO representation defined by a translationally invariant MPO tensor $W_{\alpha\beta}^{ii'}$ that is lower triangular in the bond space, this same formalism will enable us to extract the energy density of an iMPS in the thermodynamic limit, as described in Ref. [62,63]. Several works, e.g., Ref. [64], have discussed how to put common Hamiltonian terms into this lower triangular ‘‘canonical form’’ and compose them. For concreteness, we will consider the MPO representation of the TFIM used in the main text, given as

$$\hat{W} = \begin{pmatrix} \hat{I} & 0 & 0 \\ \hat{\sigma}_z & 0 & 0 \\ -h\hat{\sigma}^x & \hat{\sigma}_z & \hat{I} \end{pmatrix}. \quad (\text{A4})$$

For each value of the bond index of the MPO matrix we can associate a matrix \mathbb{E}_α , and the fixed-point relations for these matrices associated with the left action of generalized transfer operators take the form

$$\mathbb{E}_\alpha = T_{L\hat{W}_{\alpha\alpha}}(A, \mathbb{E}_\alpha) + \sum_{\beta>\alpha} T_{L\hat{W}_{\beta\alpha}}(A, \mathbb{E}_\beta). \quad (\text{A5})$$

Because of the lower triangular nature of the MPO, we can solve for the terms in order, starting from the last ($\alpha = \chi_W$, with χ_W the bond dimension of the MPO). Namely, we find the fixed-point relation

$$\mathbb{E}_{\chi_W} = T_{L\hat{I}}(A, \mathbb{E}_{\chi_W}), \quad (\text{A6})$$

Algorithm 3 Algorithm for the energy density of the TFIM with respect to a left-canonical iMPS.

Require: A in left canonical form, $\sum_i \mathbb{A}^{i\dagger} \mathbb{A}^i = \mathbb{I}$
function ENERGYDENSITY(A, h)
 $T_{(\alpha\alpha'),(\beta\beta')} \leftarrow \sum_i A_{\alpha\beta}^{i*} A_{\alpha'\beta'}^i$
 $\rho_{\text{half-infinite}} \leftarrow$ right eigenvector of T with eigenvalue 1
 $\mathbb{E}_z \leftarrow \sum_{ii'} \langle i | \hat{\sigma}_z | i' \rangle \mathbb{A}^{i\dagger} \mathbb{A}^{i'}$
 $\mathbb{E}_x \leftarrow \sum_{ii'} \langle i | \hat{\sigma}_x | i' \rangle \mathbb{A}^{i\dagger} \mathbb{E}_z \mathbb{A}^{i'}$
 $\mathbb{E}_x \leftarrow \sum_{ii'} \langle i | \hat{\sigma}_x | i' \rangle \mathbb{A}^{i\dagger} \mathbb{A}^{i'}$
 $e_0 \leftarrow \text{Tr}[\rho_{\text{half-infinite}} (\mathbb{E}_z - h\mathbb{E}_x)]$
return e_0
end function

and so $\mathbb{E}_{\chi_W} = \mathbb{I}$ for an iMPS in left-canonical form. For the case of the TFIM MPO, Eq. (A4), the next matrix is

$$\mathbb{E}_2 = T_{L\hat{\sigma}_z}(A, \mathbb{E}_{\chi_W}), \quad (\text{A7})$$

$$= T_{L\hat{\sigma}_z}(A, \mathbb{I}). \quad (\text{A8})$$

Hence, this is not a system of equations but a definition. The one remaining fixed-point relation is

$$\mathbb{E}_1 = T_{L\hat{I}}(A, \mathbb{E}_1) + T_{L\hat{\sigma}_z}(A, \mathbb{E}_2) - hT_{L\hat{\sigma}_x}(A, \mathbb{E}_{\chi_W}), \quad (\text{A9})$$

$$\Rightarrow \mathbb{E}_1(\mathbb{I} - T_I(A)) = T_{L\hat{\sigma}_z}(A, T_{L\hat{\sigma}_z}(A, \mathbb{I})) - hT_{L\hat{\sigma}_x}(A, \mathbb{I}). \quad (\text{A10})$$

Recalling that the transfer operator $T_I(A)$ has the identity as a left eigenmatrix, we see that the left-hand side of this (left) system of equations is singular, meaning that we can add any multiple of the identity to a solution and obtain another solution. We can find the solution with zero weight on the identity by removing the components in this direction from the right-hand side matrix. This is achieved by recalling that $\rho_{\text{half-infinite}}$ is a right eigenmatrix of the transfer operator with eigenvalue 1, and so spans the null space of the left-hand operator. Hence, subtracting \mathbb{I} times

$$e_0 = \text{Tr}[\rho_{\text{half-infinite}} T_{L\hat{\sigma}_z}(A, T_{L\hat{\sigma}_z}(A, \mathbb{I})) - hT_{L\hat{\sigma}_x}(A, \mathbb{I})], \quad (\text{A11})$$

from the right-hand side ensures that the right-hand side is orthogonal to the null space of the left-hand side, and so this equation has a consistent solution in spite of being singular. The extracted quantity e_0 is the energy density. The singularity in the relation for \mathbb{E}_1 amounts to an ambiguity in defining the zero of energy.

The above observations lead directly to an algorithm for the energy density of a left-canonical iMPS in the thermodynamic limit, which for the TFIM takes the form of Algorithm 3.

This algorithm can hence be used as a cost function in a minimization routine to optimize the tensor A . In order to ensure that the tensor remains in left-canonical form during the optimization, we utilize Riemannian manifold gradient descent with the matrix $\tilde{A}_{(\alpha i)\beta} = A_{\alpha\beta}^i$ forming an element of the Stiefel manifold $\text{St}(\chi d, \chi)$ [65]. As detailed in other recent works [35,48,49], this procedure consists of projection of cost function gradients onto the tangent space of the manifold at the current estimate point and subsequent update of the tensor using a retraction of this projected gradient.

Algorithm 4 Cost function for difference between true half infinite density matrix and that obtained from a finite burn-in procedure

Require: A in left canonical form, $\sum_i \mathbb{A}^{i\dagger} \mathbb{A}^i = \mathbb{I}$
Require: V with unit norm, $\sum_\alpha |V_\alpha|^2 = 1$
function BURNINCOST(A, V, N'_b)
 $T_{(\alpha\alpha'),(\beta\beta')} \leftarrow \sum_i A_{\alpha\beta}^{i*} A_{\alpha'\beta'}^i$
 $\rho_{\text{half-infinite}} \leftarrow$ right eigenvector of T with eigenvalue 1
 $B_\alpha^i \leftarrow \sum_\beta A_{\alpha\beta}^i V_\beta$
 $U_{\alpha\alpha'} \leftarrow \sum_i B_{\alpha'}^{i*} B_\alpha^i$
 $N \leftarrow 1$
while $N < N'_b$ **do**
 $Q_{\alpha\beta}^i \leftarrow \sum_{\alpha'} U_{\alpha\alpha'} A_{\alpha'\beta}^i$
 $U_{\alpha\alpha'} \leftarrow \sum_{\beta,i} Q_{\beta\alpha'}^i A_{\beta\alpha}^{i*}$
 $N \leftarrow N + 1$
end while
return norm($U - \rho_{\text{half-infinite}}$)
end function

Vector transport can be used to incorporate information from previous iterations through the momentum technique. This process integrates readily with modern numerical packages performing automatic differentiation to obtain gradients, such as TensorFlow [66] and PyTorch [53].

The tensor output from this optimization procedure produces the proper half-infinite density matrix as the fixed point of the transfer matrix. However, when practically implemented on quantum hardware we will produce an approximation to the half-infinite density matrix by using the ‘‘burn in’’ procedure with a fixed number N'_b of iterations (taken as a hyperparameter in our approach). Hence, we will optimize the initial state of the bond qubits (boundary vector) to reproduce the true half-infinite density matrix obtained as the fixed point of the transfer operator as closely as possible using this procedure. The cost function for this process is given as Algorithm 4.

Algorithm 5 Online algorithm to determine class from probabilistic classifier with confidence parameter p_*

Require: p_* , method to obtain probabilistic class prediction
SamplePrediction
function CLASSFROMSAMPLES(p_*)
 $B \leftarrow 4, n_0 \leftarrow 0, n_1 \leftarrow 0, p < \leftarrow \frac{1}{2}$
while **True** **do**
 $s \leftarrow$ **SamplePrediction**
if $s == 0$ **then**
 $p < \leftarrow p < - \frac{1}{(n_0+1)B}$
 $B \leftarrow 2B \frac{n_0+1}{n_0+n_1+2}$
 $n_0 \leftarrow n_0 + 1$
else
 $p < \leftarrow p < + \frac{1}{(n_1+1)B}$
 $B \leftarrow 2B \frac{n_1+1}{n_0+n_1+2}$
 $n_1 \leftarrow n_1 + 1$
end if
if $p < > p_*$ **then return** 1
else if $p < < 1 - p_*$ **then return** 0
end if
end while
end function

This cost function can again be integrated with manifold optimization on the χ -dimensional sphere manifold to optimize the boundary state V . With the optimized tensor in hand, we can utilize the gauging and compilation procedures described in the main text to convert this into a quantum circuit for implementation on hardware.

APPENDIX B: ALGORITHM FOR SAMPLING A CLASSICAL PRODUCT STATE FROM A MATRIX PRODUCT STATE

In this Appendix, we review an efficient algorithm for sampling a classical product state in the basis defined by a collection of L Hermitian operators $\{\hat{O}_i, i = 0, \dots, L-1\}$ from a finite MPS $|\psi[A]\rangle$ on L sites with open boundary conditions ($\langle l| = |r\rangle = \delta_{\alpha,0}$) according to its probability of occurrence in the MPS. Specifically, product states $|\lambda_0, \dots, \lambda_{L-1}\rangle$ in a basis to be described will be sampled from $|\langle \psi[A] | \lambda_0, \dots, \lambda_{L-1} \rangle|^2$. This algorithm first appeared in the context of minimally entangled typical thermal state (METTS) algorithms [67]. Given the Hermitian operators \hat{O}_i with spectral decomposition $\hat{O}_i = \sum_{\lambda_i} \lambda_i |\lambda_i\rangle \langle \lambda_i|$ we can define projection operators $\hat{P}_{i,\mu} = |\lambda_{i,\mu}\rangle \langle \lambda_{i,\mu}|$. Now, assuming the MPS is in left-canonical form, we sample the $(L-1)$ th site by sampling from the categorical distribution defined by the vector $v_\mu = \text{Tr}[\hat{P}_\mu \hat{\rho}_{L-1}]$, in which $[\hat{\rho}_{L-1}]_{j,j'} = \sum_\alpha A_{\alpha,0}^{[L-1]j} A_{\alpha,0}^{[L-1]j'*}$ is the $(L-1)$ th site’s reduced density matrix. Given the index chosen by the random sampling, μ_{L-1} , we now construct an MPS representation of $\hat{P}_{L-1,\mu_{L-1}} |\psi[A]\rangle$ by replacing $A_\alpha^{[L-1]} \leftarrow \sum_j \langle j | \lambda_{\mu_{L-1}} \rangle A_{\alpha,0}^{[L-1],j}$. Since the $(L-1)$ th site now no longer has a site index, we can absorb this tensor into the tensor for site $(L-2)$ as

$$A_{\alpha,0}^{[L-2]j} \leftarrow \sum_\beta A_{\alpha\beta}^{[L-2]j} A_\beta^{[L-1]}, \quad (\text{B1})$$

to define an MPS on $(L-1)$ sites describing $\hat{P}_{L-1,\mu_{L-1}} |\psi[A]\rangle$. This process is now repeated, constructing the single-site reduced density matrix for site $(L-2)$, sampling from the (normalized) vector obtained by tracing the reduced density matrix with the projectors onto the eigenspaces, applying the chosen projector to the state, and contracting the MPS tensor from the sampled site into the neighboring MPS tensor. Repeating for all remaining sites, the result of this algorithm is the string $(\mu_0, \dots, \mu_{L-1})$ which can be re-interpreted as a CPS $|\lambda_{\mu_0}, \dots, \lambda_{\mu_{L-1}}\rangle$. This algorithm produces CPS sampled according to $|\langle \lambda_{\mu_0}, \dots, \lambda_{\mu_{L-1}} | \psi[A] \rangle|^2$, and has a $O(\chi^2 L)$ cost that is subleading to the $O(L\chi^3)$ cost of variational ground state search.

APPENDIX C: ONLINE ALGORITHM FOR DETERMINING THE NUMBER OF SHOTS NEEDED FOR ACCURATE CLASSIFICATION

For classically preconditioned models such as those described in this paper, the Wilson score interval can be used to estimate the number of experimental shots required to infer the class of a given input state $|\mathbf{x}\rangle$ with a desired accuracy, as p_0 can be obtained from the classically tractable model. As models are scaled up towards the classically intractable regime, we no longer have knowledge of p_0 , and it must instead be inferred from the experimental data. In this scenario,

it is useful to have an *online* algorithm which estimates a distribution over p_0 conditioned on the shot history, such that statistical measures on this conditional distribution can determine when a sufficient number of shots have been taken. A natural setting for this is Bayesian inference. Using a uniform prior distribution over p_0 , corresponding to the absence of any *a priori* knowledge about p_0 , and the assumption that each shot of the experiment is a Bernoulli trial and hence described by a binomial distribution, the resulting posterior distribution for p_0 conditioned on a shot history of n_0 zeros and n_1 ones is a beta distribution

$$P(p_0|n_0, n_1) = \frac{1}{B(n_0 + 1, n_1 + 1)} p_0^{n_0} (1 - p_0)^{n_1}, \quad (\text{C1})$$

in which $B(x, y)$ is the beta function. Hence, the probability that $p_0 \leq \frac{1}{2}$, corresponding to the posterior model predicting class one, is given by the cumulative distribution function of this posterior distribution evaluated at $\frac{1}{2}$. For the present case, this is

$$p_{<} \equiv \int_0^{1/2} dp_0 P(p_0|n_0, n_1) = I_{\frac{1}{2}}(n_0 + 1, n_1 + 1), \quad (\text{C2})$$

in which $I_x(a, b)$ is the regularized incomplete beta function. With this, one can set a threshold $\frac{1}{2} < p_{\star} < 1$ such that no more shots are taken and the class is called as one when $p_{<} > p_{\star}$ or zero when $p_{<} < 1 - p_{\star}$. These criteria define that the interval $[0, 1/2]$ is a credible interval with probability p_{\star} , in the Bayesian sense.

Using the identities

$$I_x(a + 1, b) = I_x(a, b) - \frac{x^a(1-x)^b}{aB(a, b)}, \quad (\text{C3})$$

$$I_x(a, b + 1) = I_x(a, b) - \frac{x^a(1-x)^b}{bB(a, b)}, \quad (\text{C4})$$

$$B(a + 1, b) = B(a, b) \frac{a}{a + b}, \quad (\text{C5})$$

$$B(a, b + 1) = B(a, b) \frac{b}{a + b}, \quad (\text{C6})$$

and $I_{\frac{1}{2}}(1, 1) = \frac{1}{2}$, $B(1, 1) = 1$, we can develop an online Bayesian algorithm that will halt when the requisite number of shots have been taken, summarized in Algorithm 5.

From this algorithm, we can also obtain the maximum *a posteriori* estimate of p_0 as $\frac{n_0}{n_0 + n_1}$.

-
- [1] A. W. Harrow and A. Montanaro, *Nature (London)* **549**, 203 (2017).
- [2] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell *et al.*, *Nature (London)* **574**, 505 (2019).
- [3] H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu *et al.*, *Science* **370**, 1460 (2020).
- [4] J. I. Cirac and P. Zoller, *Nat. Phys.* **8**, 264 (2012).
- [5] N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn *et al.*, *Quantum Sci. Technol.* **3**, 030503 (2018).
- [6] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature (London)* **549**, 195 (2017).
- [7] A. Montanaro, *npj Quantum Inf.* **2**, 15023 (2016).
- [8] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, *Phys. Rev. A* **77**, 012307 (2008).
- [9] W. Rawat and Z. Wang, *Neural Comput.* **29**, 2352 (2017).
- [10] A. Sherstinsky, *Physica D* **404**, 132306 (2020).
- [11] U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).
- [12] A. Cichocki, *arXiv:1407.3124*.
- [13] R. Orús, *Ann. Phys.* **349**, 117 (2014).
- [14] G. K.-L. Chan, A. Keselman, N. Nakatani, Z. Li, and S. R. White, *J. Chem. Phys.* **145**, 014102 (2016).
- [15] R. Orús, *Nat. Rev. Phys.* **1**, 538 (2019).
- [16] J. M. Dudek, L. Duenas-Osorio, and M. Y. Vardi, *arXiv:1908.04381*.
- [17] G. Torlai, C. J. Wood, A. Acharya, G. Carleo, J. Carrasquilla, and L. Aolita, *arXiv:2006.02424*.
- [18] M. M. Wolf, F. Verstraete, M. B. Hastings, and J. I. Cirac, *Phys. Rev. Lett.* **100**, 070502 (2008).
- [19] J. Eisert, *Model. Simulat.* **3**, 520 (2013).
- [20] S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [21] A. van den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, *arXiv:1609.03499*.
- [22] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, *Quantum Sci. Technol.* **4**, 024001 (2019).
- [23] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, *Phys. Rev. Lett.* **95**, 110503 (2005).
- [24] T. J. Osborne, J. Eisert, and F. Verstraete, *Phys. Rev. Lett.* **105**, 260401 (2010).
- [25] J. I. Cirac, D. Poilblanc, N. Schuch, and F. Verstraete, *Phys. Rev. B* **83**, 245134 (2011).
- [26] I. H. Kim, *arXiv:1702.02093*.
- [27] I. H. Kim, *arXiv:1703.00032*.
- [28] I. H. Kim and B. Swingle, *arXiv:1711.07500*.
- [29] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter, *Phys. Rev. Research* **3**, 033002 (2021).
- [30] J. Yirka and Y. Subaşı, *Quantum* **5**, 535 (2021).
- [31] M. Foss-Feig, S. Ragole, A. Potter, J. Dreiling, C. Figgatt, J. Gaebler, A. Hall, S. Moses, J. Pino, B. Spaun *et al.*, *Phys. Rev. Lett.* **128**, 150504 (2022).
- [32] E. Chertkov, J. Bohnet, D. Francois, J. Gaebler, D. Gresh, A. Hankin, K. Lee, R. Tobey, D. Hayes, B. Neyenhuis *et al.*, *arXiv:2105.09324*.
- [33] Z. Liu, L.-W. Yu, L.-M. Duan, and D.-L. Deng, *arXiv:2108.08312*.
- [34] M. L. Wall, M. R. Abernathy, and G. Quiroz, *Phys. Rev. Research* **3**, 023010 (2021).
- [35] M. L. Wall and G. D'Aguanno, *Phys. Rev. A* **104**, 042408 (2021).
- [36] A. V. Uvarov, A. S. Kardashin, and J. D. Biamonte, *Phys. Rev. A* **102**, 012415 (2020).
- [37] W. S. Bakr, J. I. Gillen, A. Peng, S. Fölling, and M. Greiner, *Nature (London)* **462**, 74 (2009).

[38] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas, *Phys. Rev. Lett.* **100**, 200502 (2008).

[39] L. A. Zhukas, P. Svihra, A. Nomerotski, and B. B. Blinov, *Phys. Rev. A* **103**, 062614 (2021).

[40] M. B. Hastings, *Phys. Rev. B* **73**, 085115 (2006).

[41] J. Pino, J. Dreiling, C. Figgatt, J. Gaebler, S. Moses, M. Allman, C. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer *et al.*, *Nature (London)* **592**, 209 (2021).

[42] J. Gaebler, C. Baldwin, S. Moses, J. Dreiling, C. Figgatt, M. Foss-Feig, D. Hayes, and J. Pino, *Phys. Rev. A* **104**, 062440 (2021).

[43] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, *Phys. Rev. X* **8**, 031012 (2018).

[44] E. Stoudenmire and D. J. Schwab, in *Proceedings of the 30th International Conference on Neural Information Processing Systems, Barcelona, 2016*, edited by D. D. Lee, U. von Luxburg, R. Garnett, M. Sugiyama, and I. Guyon (Curran, Red Hook, 2016), pp. 4799–4807.

[45] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Nat. Commun.* **9**, 4812 (2018).

[46] E. Grant, L. Wossnig, M. Ostaszewski, and M. Benedetti, *Quantum* **3**, 214 (2019).

[47] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, *Nat. Commun.* **12**, 1791 (2021).

[48] I. Luchnikov, M. Krechetov, and S. Filippov, *New J. Phys.* **23**, 073006 (2021).

[49] M. Hauru, M. Van Damme, and J. Haegeman, *Sci. Post Phys.* **10**, 040 (2021).

[50] N. Boumal, B. Mishra, P.-A. Absil, and R. Sepulchre, *J. Machine Learn. Res.* **15**, 1455 (2014).

[51] J. Townsend, N. Koep, and S. Weichwald, *J. Machine Learn. Res.* **17**, 4755 (2016).

[52] J. Preskill, Lecture Notes for PH219/CS219: Quantum Information and Computation, Caltech, http://theory.caltech.edu/~preskill/ph219/ph219_2020-21.html.

[53] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Vancouver, 2019), pp. 8024–8035.

[54] S. Sachdev, *Quantum Phase Transitions* (Cambridge University Press, New York, 2011).

[55] K. Kim, S. Korenblit, R. Islam, E. Edwards, M. Chang, C. Noh, H. Carmichael, G. Lin, L. Duan, C. J. Wang *et al.*, *New J. Phys.* **13**, 105003 (2011).

[56] R. Toskovic, R. Van Den Berg, A. Spinelli, I. Eliens, B. Van Den Toorn, B. Bryant, J.-S. Caux, and A. Otte, *Nat. Phys.* **12**, 656 (2016).

[57] A. Dutta, G. Aeppli, B. K. Chakrabarti, U. Divakaran, T. F. Rosenbaum, and D. Sen, Experimental realizations of transverse field ising systems, in *Quantum Phase Transitions in Transverse Field Spin Models: From Statistical Physics to Quantum Information* (Cambridge University Press, Delhi, 2015), p. 231–244.

[58] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg *et al.*, *J. Machine Learn. Res.* **12**, 2825 (2011).

[59] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/PhysRevA.105.062439> for quantum assembly language (QASM) representations of all quantum circuits run on the H1 machine.

[60] A. J. Ferris and D. Poulin, *Phys. Rev. Lett.* **113**, 030501 (2014).

[61] A. Jahn and J. Eisert, *Quantum Sci. Technol.* **6**, 033002 (2021).

[62] I. P. McCulloch, [arXiv:0804.2509](https://arxiv.org/abs/0804.2509).

[63] L. Michel and I. P. McCulloch, [arXiv:1008.4667](https://arxiv.org/abs/1008.4667).

[64] I. P. McCulloch, *J. Stat. Mech.: Theory Exp.* (2007) P10014.

[65] A. Edelman, T. A. Arias, and S. T. Smith, *SIAM J. Matrix Anal. Appl.* **20**, 303 (1998).

[66] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard *et al.*, in *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)* (USENIX association, Savannah, GA, 2016), pp. 265–283.

[67] E. Stoudenmire and S. R. White, *New J. Phys.* **12**, 055026 (2010).