Generative machine learning with tensor networks: Benchmarks on near-term quantum computers

Michael L. Wall, Matthew R. Abernathy , and Gregory Quiroz

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

(Received 16 October 2020; revised 3 March 2021; accepted 7 March 2021; published 2 April 2021)

Noisy, intermediate-scale quantum (NISQ) computing devices have become an industrial reality in the last few years, and cloud-based interfaces to these devices are enabling the exploration of near-term quantum computing on a range of problems. As NISQ devices are too noisy for many of the algorithms with a known quantum advantage, discovering impactful applications for near-term devices is the subject of intense research interest. We explore a quantum-assisted machine learning (QAML) workflow using NISQ devices through the perspective of tensor networks (TNs), which offer a robust platform for designing resource-efficient and expressive machine learning models to be dispatched on quantum devices. In particular, we lay out a framework for designing and optimizing TN-based QAML models using classical techniques, and then compiling these models to be run on quantum hardware, with demonstrations for generative matrix product state (MPS) models. We put forth a generalized canonical form for MPS models that aids in compilation to quantum devices, and demonstrate greedy heuristics for compiling with a given topology and gate set that outperforms known generic methods in terms of the number of entangling gates, e.g., CNOTs, in some cases by an order of magnitude. We present an exactly solvable benchmark problem for assessing the performance of MPS QAML models and also present an application for the canonical MNIST handwritten digit dataset. The impacts of hardware topology and day-to-day experimental noise fluctuations on model performance are explored by analyzing both raw experimental counts and statistical divergences of inferred distributions. We also present parametric studies of depolarization and readout noise impacts on model performance using hardware simulators.

DOI: [10.1103/PhysRevResearch.3.023010](https://doi.org/10.1103/PhysRevResearch.3.023010)

I. INTRODUCTION

In recent years, gate-based quantum computing has emerged as a relatively mature technology, with many platforms offering cloud-based interfaces to machines with a few to dozens of qubits [1–5], as well as classical emulators of quantum devices of this class [6]. Today’s quantum computing resources remain a long way from the millions of qubits [7] required to perform canonical quantum computing tasks such as integer factorization with error correction [8,9], and present devices are either engineered with a specific demonstration goal or designed for general-purpose research-scale exploration [10]. With the advent of noisy, intermediate-scale quantum (NISQ) devices [11], whose hardware noise and limited qubit connectivity and gate sets pose challenges for demonstrating scalable universal quantum computation, we are faced with a different form of quantum application discovery in which algorithms need to be robust to noise, limited qubit connectivity and gate sets, and highly resource-efficient.

Machine learning (ML) has been put forward as a possible application area for NISQ devices, with a range of recent proposals [12–14]. ML may prove promising for NISQ

applications because well-performing ML algorithms feature robustness against noise, quantum circuits can be designed for ML applications that are highly qubit-efficient [15], and quantum models can be designed whose expressibility increases exponentially with qubit depth [15,16]. The most impactful near-term ML application likely lies in quantum-assisted machine learning (QAML), in which a quantum circuit’s parameters are classically optimized based on measurement outcomes that may not be efficiently classically simulable [17]; this also includes kernel-based learning schemes with a quantum kernel [18]. Tensor networks (TNs) provide a robust means of designing such parameterized quantum circuits that are quantum-resource efficient and can be implemented and optimized on classical or quantum hardware. TN-based QAML algorithms hence leverage the significant research effort into optimization strategies for TNs [19–21], and also enable detailed benchmarking and design of QAML models classically, with a smooth transition to classically intractable models.

In this work, we explore the applicability of QAML with TN architectures on NISQ hardware and hardware simulators, exploring the effects of present-day and near-term hardware noise, qubit connectivity, and restrictions on gate sets. While previous studies have investigated the impacts of noise on ML model outcomes [15,22], our exposition significantly increases the fidelity of such analyses by obtaining an explicit representation of the optimal TN model with fixed resources for a given target quantum hardware. Also in contrast to previous hardware demonstrations on a classification task [22], we

FIG. 1. Overview of QAML workflow. Classical data in (a) is preprocessed and transformed to quantum states embedded in an exponentially large Hilbert space in (b). A TN model is learned from a collection of quantum training data in (c), which has the interpretation in (d) of a sequential preparation scheme involving a small number of readout qubits coupled to ancillary resources. The isometries of the sequential preparation scheme in (d) are conditioned using inherent freedom in the TN representation in (e), and then converted into native gates for a target hardware architecture, displayed as the IBMQ-X2 processor for concreteness. Running on cloud-based hardware in (f), we obtain measurements defining output predictions, as in (g). For interpretation of graphical representations, see text.

focus on fully generative unsupervised learning tasks, which have been identified as a promising avenue for QAML [13], and focus on the most resource-efficient matrix product state (MPS) TN topology. We present a framework for QAML—outlined in Fig. 1—that includes translation of classical data into quantum states, optimization of an MPS model using classical techniques, the conversion of this classically trained model into a sequence of isometric operations to be performed on quantum resources, and the optimization and compilation of these isometric operations into native operations for a given hardware topology and allowed gate set. In particular, we develop several novel techniques for the compilation stage aimed at TN models for QAML on NISQ devices, such as the permutation of auxiliary quantum degrees of freedom in the TN to optimize mapping to hardware resources and heuristics for the translation of isometries into native operations using as few entangling operations (e.g., CNOTs) as possible.

The tools developed herein enable the robust design and performance assessment of QAML models on NISQ devices in the regime where classical simulations are still possible, and will inform architectures and noise levels for scaling to the classically intractable regime. Even in the classically intractable regime in which the model must be optimized using a quantum device in a hybrid quantum/classical loop [23,24], our techniques provide a means of obtaining an approximate,

classically trained “preconditioner” for the quantum models that can help avoid local minima and reduce optimization time, as has been explored using classical simulations for other architectures in Ref. [22]. We present exemplar results for our workflow for synthetic data that can be described by an exactly solvable two-qubit MPS QAML model, as well as on features extracted from the canonical MNIST handwritten digit dataset [25].

The remainder of this work is organized as follows: Sec. II discusses QAML with tensor networks (TNs) broadly, including embedding of classical data into quantum states, classical training of a TN model, and the conversion of TN models into resource-efficient sequential preparation schemes; Sec. III discusses our approach for compiling TN-based QAML models for running on quantum hardware, including the utilization of ambiguity in the TN representation and greedy compilation heuristics for minimizing model gate depth; Sec. IV presents an exactly solvable two-qubit QAML model and assesses the performance of our QAML workflow on quantum hardware; in Sec. V, we give an example application to generative modeling of features extracted from the MNIST dataset and analyze the performance of our models as a function of hardware noise using a quantum hardware simulator; finally, in Sec. VI, we conclude and give an outlook. Details of the MNIST model studied in Sec. V are given in Appendix E.

II. QUANTUM-ASSISTED MACHINE LEARNING WITH TENSOR NETWORKS

Figure 1 broadly outlines the QAML workflow explored in the present work. We begin with a collection of classical data vectors in a training set $\mathcal{T} = \{\mathbf{x}_j\}_{j=1}^{N_r}$, where each element $\mathbf{x}_j$ is an N -length vector. The first step in our QAML workflow is to define a mapping of classical data vectors to vectors in a quantum Hilbert space. Here, the only restriction we will place on the encoding of classical data in quantum states is that each classical data vector is encoded in an unentangled product state. This is useful for several reasons. For one, unentangled states are the simplest to prepare experimentally with high fidelity, and also enable us to use qubit-efficient sequential preparation schemes. From a learning perspective, encoding individual data vectors in product states ensures that any entanglement that results in a quantum model comes from correlations in an ensemble of data and not from *a priori* assumptions about pre-existing correlations for individual data vectors [26].

The simplest case of a quantum embedding occurs when the data is discrete, and so can be formulated as vectors $\mathbf{x}$ whose individual elements $x_i \in \{0, 1\}$. In this case each element is mapped to a qubit as $|x_i\rangle$, such that the embedding of the full N -dimensional classical data vector into a register of N qubits is [27,28]

$$|\mathbf{x}\rangle = \bigotimes_{i=1}^{N-1} |x_i\rangle. \quad (1)$$

Embeddings can also be formulated for vectors of continuous variable data, as has been explored in Refs. [26,29–31]. We will not review those cases here, as only the binary embedding Eq. (1) is utilized in this work.

A. Tensor networks and sequential preparation

The next step in our QAML workflow outlined in Fig. 1 is to learn a quantum model for the collection of quantum states $\{|\mathbf{x}_j\rangle\}_{j=1}^{N_r}$ resulting from applying the encoding map from the previous section to the training data. Here, we define a quantum model as a collection of operations applied to quantum resources to produce a state that encodes the properties of the ensemble $\{|\mathbf{x}_j\rangle\}$. In what follows, we specialize to the case of tensor network (TN) models, which provide a convenient parametrization of the structure of quantum operations and resources. Generally speaking, TNs represent the high-rank tensor describing a quantum wave function in a specified basis as a contraction over low-rank tensors, and hence define families of low-rank approximations whose computational power can be expressed in terms of the maximum dimension of any contracted index χ , known as the bond dimension.

A wide variety of TN topologies have been considered which are able to efficiently capture certain classes of quantum states [19–21]; in the present work, we focus on matrix product states (MPSs). MPSs use a one-dimensional TN topology, as shown using the Penrose graphical notation for tensors [19] in Fig. 1(c), and form the basis for the enormously successful density matrix renormalization group (DMRG) algorithm in quantum condensed matter physics [32]. MPSs have several properties that make them attractive for QAML.

For one, they are undoubtedly the most well-understood and mature of all tensor networks, which has led to robust optimization strategies that are widely used in the quantum many-body community. In addition, MPSs are highly quantum resource efficient, in the sense that their associated wave functions can be sequentially prepared, and so qubits can be reused in deployment on quantum hardware. In fact, it can be shown that every state that can be sequentially prepared can be written as an MPS [33–35]. This last point importantly means that any QAML model that achieves the optimal $\mathcal{O}(1)$ scaling of qubit requirements with data vector length N using sequential preparation can be expressed as an MPS, and the methods and analyses presented in this work apply.

In recent years, TNs have found applications outside of the condensed matter and quantum information domains. The mathematical analysis community has proposed TN methods for data analysis, e.g., large-scale principal component analysis [36,37]. In this community, MPSs are referred to as tensor trains [38]. Using TN methods to design quantum-inspired ML models was first proposed by Stoudenmire and Schwab [26], who put forth a scheme using a MPS network as a linear classifier in a Hilbert space whose dimension is exponentially large in the length of the raw data vector. Since then, many other proposals for quantum-assisted or quantum-inspired TN ML models have appeared in the literature [22,39–56], including generative modeling of binary data using MPSs in Ref. [57]. In the majority of approaches, DMRG-inspired algorithms for optimization have been employed. However the authors of Ref. [58] recently demonstrated an alternate strategy where a TN was implemented as a neural network using standard deep learning software, and the tensors of the TN were optimized using backpropagation strategies ubiquitous in classical ML. While this strategy has shown good performance, it has also been shown to be suboptimal with respect to the DMRG-like approach [59]. Nonetheless, the use of deep learning “preconditioners” and the intersection of QAML and neural networks remains intriguing [60–62].

The fact that MPSs define a sequential preparation scheme means that MPSs provide highly resource efficient architectures for learning [15] and quantum simulation [63]. In particular, the qubit resource requirements for an MPS model are logarithmic in the bond dimension χ , which encapsulates the expressivity of the model, and are independent of the length of the input data vector N . The details of the procedure connecting MPS models and sequential preparation schemes are reviewed in Appendix A, where it is shown that a single “physical” or “readout” qubit together with a χ -level ancilla (which itself can be made of multiple qubits) suffices to extract information from an MPS model, provided we are not measuring any multiqubit properties of the state. As shown in Appendix A, a sequential preparation procedure defined as an MPS takes the form

$$|\psi\rangle = \sum_{j_0 \dots j_{N-1}} \text{Tr}[\mathbb{L}^{[0]j_0} \dots \mathbb{L}^{[N-1]j_{N-1}}] |j_0 \dots j_{N-1}\rangle, \quad (2)$$

in which the $\chi \times \chi$ matrices $\mathbb{L}^{[i]j_i}$ satisfy the isometry condition

$$\sum_{j_i} \mathbb{L}^{[i]j_i \dagger} \mathbb{L}^{[i]j_i} = \mathbb{I}_\chi. \quad (3)$$

Those familiar with the theory of MPSs will recognize Eq. (3) as defining the left-canonical form of an MPS; standard procedures exist for putting an MPS into this form [19–21].

The condition of not measuring any multiqubit properties of the state holds for our specified use case of generating data vector samples a single element at a time, as shown in Fig. 1 (d). Here, the single physical qubit is coupled to three ancilla qubits forming a $\chi = 8$ -level resource, and the physical qubit is measured, with the measurement outcome forming the data sample element x_{N-1} . The ancilla qubits are left unmeasured, the physical qubit is reinitialized, and the procedure of coupling the physical and ancilla qubits and measuring the latter repeated, resulting in the remaining elements of the data vector $x_{N-2}, \dots, x_0$. The process of coupling the physical $\{|j_i\rangle\}$ and ancilla $\{|\alpha\rangle\}$ states is defined by the isometric operators

$$\hat{L}^{[i]} = \sum_{\alpha\beta j_i} L_{\alpha\beta}^{[i]j_i} |\alpha j_i\rangle \langle \beta 0|. \quad (4)$$

The sequential preparation scheme with a single readout qubit requires mid-circuit measurement and reset which is not universally available in present-day hardware, but has been demonstrated in, e.g., trapped ion platforms [64]. Finally, we stress that the scheme for converting from an MPS model to a sequential preparation procedure is formal in the sense that it produces isometries acting on quantum resources without reference to their actual physical representation or other hardware constraints such as limited coherence time, connectivity, gate sets, etc.. The translation of these formal isometries into operations to be dispatched on a given target hardware are detailed in Sec. III.

B. Generative MPS models and classical training procedure

We now further specialize to generative models, in which a collection of quantum data vectors are encoded into a wave function $|\psi\rangle$ such that the probability distribution evaluated at data vector $\mathbf{x}$ is

$$P(\mathbf{x}) = \frac{\langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle}{Z}. \quad (5)$$

Here, $Z = \langle \psi | \psi \rangle = \sum_{\mathbf{x}} \langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle$ is a normalization factor, and $|\mathbf{x}\rangle$ denotes the encoding of a classical data vector $\mathbf{x}$ into a quantum state as in Eq. (1). As this corresponds to Born's rule for measurement outcomes, the resulting structure is referred to as a *Born machine* [57,65]. The state $|\psi\rangle$, or, equivalently, the process of generating $|\psi\rangle$ from a known fiducial state, is our quantum machine learning model.

In order to discuss data representation using Born machines, we define the average log-likelihood of the data in the training set $\mathcal{T}$ as

$$\mathcal{L}(\mathcal{T}) = \frac{1}{N_T} \sum_{\mathbf{x} \in \mathcal{T}} \ln \left[\frac{\langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle}{Z} \right]. \quad (6)$$

The minimization of the negative log-likelihood with respect to the parameters in our Born machine is equivalent to maximizing the probability that the data is generated by the Born machine. Parameterizing the wave function $|\psi\rangle$ as an MPS as

$$|\psi\rangle = \sum_{i_0 \dots i_{N-1}} \text{Tr}[\mathbb{A}^{i_0^\dagger} \dots \mathbb{A}^{i_{N-1}^\dagger}] |i_0 \dots i_{N-1}\rangle, \quad (7)$$

and using the binary encoding of data, Eq. (1), we find

$$\mathcal{L}(\mathcal{T}) = \frac{1}{N_T} \sum_{\mathbf{x} \in \mathcal{T}} \ln \left[\frac{1}{Z} |\text{Tr}[\mathbb{A}^{x_0^\dagger} \dots \mathbb{A}^{x_{N-1}^\dagger}]|^2 \right],$$

where the normalization factor (partition function) is

$$Z = \sum_{i_0 \dots i_{N-1}} \text{Tr}[\mathbb{A}^{i_0^\dagger} \dots \mathbb{A}^{i_{N-1}^\dagger}] \text{Tr}[\mathbb{A}^{i_0} \dots \mathbb{A}^{i_{N-1}}]. \quad (8)$$

We will optimize the Born machine by a DMRG-style procedure using gradient descent, where the gradient is taken with respect to the tensors of the MPS. Namely, we will consider the gradient with respect to a group of s neighboring tensors $\Theta = \mathbb{A}^{i_1} \dots \mathbb{A}^{i_{1+s}}$, with s typically being one or two, noting that the gradient of an object with respect to a tensor is a tensor whose elements are the partial derivatives with respect to the individual tensor elements. We take the gradient with respect to the conjugates of the tensors $\mathbb{A}^{i_j}$, formally considering these conjugates independent of the tensors themselves. This gradient may be written as

$$\begin{aligned} \nabla_{\Theta^\star} \mathcal{L}(\mathcal{T}) &= \frac{1}{N_T} \sum_{\mathbf{x} \in \mathcal{T}} \frac{\nabla_{\Theta^\star} \langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle}{\langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle} - \frac{\nabla_{\Theta^\star} Z}{Z} \\ &= \frac{1}{N_T} \sum_{\mathbf{x} \in \mathcal{T}} \frac{\nabla_{\Theta^\star} \langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle}{\langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle} - \sum_{\mathbf{x}} \frac{\nabla_{\Theta^\star} \langle \psi | \mathbf{x} \rangle \langle \mathbf{x} | \psi \rangle}{Z}. \end{aligned} \quad (9)$$

With this gradient in hand, we update the local block of tensors as

$$\Theta \rightarrow \Theta + \eta \nabla_{\Theta^\star} \mathcal{L}(\mathcal{T}), \quad (10)$$

in which η is a learning rate (note that this is equivalent to minimizing the negative log likelihood). For the single-site algorithm ($s = 1$), this update does not change the bond dimension or canonical form of the MPS. For the two-site algorithm ($s = 2$), we can now split the updated tensor Θ into its component MPS tensors as

$$\Theta_{\alpha\beta}^{ij} = \sum_{\mu} A_{\alpha\mu}^i A_{\mu\beta}^j, \quad (11)$$

using, e.g., the SVD. Hence, the addition of the gradient can increase the bond dimension, and thus the representation power, adaptively based on the data. The bond dimension can also be set implicitly by requiring that the L_2 -norm of the tensor Θ is represented with a bounded relative error ε . The above update has affected only a small group of the tensors with all others held fixed. A complete optimization cycle, or “sweep,” occurs when we have updated all tensors twice, moving in a back-and-forth motion over the MPS. The sweeping process is converged once the negative log-likelihood no longer decreases substantially. Example convergence behavior will be given later in Sec. V.

III. COMPILATION OF MPS MODELS FOR QUANTUM HARDWARE

In this section, we address how to take an MPS model resulting from the classical optimization procedure outlined in Sec. II B and convert it into a sequence of operations to be

performed on a quantum device. We will refer to this operation as *quantum compilation*. Many modern NISQ software ecosystems and other stand-alone applications [1,2,4,66–68] have routines for compiling quantum instructions, usually required to be supplied in the form of an abstract quantum circuit model. These compilers typically perform multiple passes through the abstract circuit to map virtual qubits from the abstract model onto the hardware qubits of the device, route operations between the virtual qubits to hardware qubits, e.g., by placing SWAP gates, and optimization to minimize some property of the circuit, such as entangling gate count. We note that quantum compilation remains an active area of research, and currently available *generic* methods for quantum compilation tend to produce “deep” circuits with significant numbers of entangling gates.

There are several unique properties of our particular quantum computing use case—compiling the isometries encoding TN models for QAML—that make them unique compared to traditional quantum computing use cases. For one, the isometric operations are defined on the Hilbert space of a physical qubit and a formal χ -level ancilla, and so may not uniquely describe an isometric operation on a set of virtual qubits, e.g., when χ is not a power of 2. Further, since the ancilla degrees of freedom are never directly measured, there is no preferred basis or state ordering for these states. Both of these properties give freedom that can be utilized to simplify compilation. In addition, the isometries are the result of an optimization procedure that has a finite tolerance (see Sec. II B), and so do not need to be compiled exactly to meet some fine-tuned property. That is to say, model predictions are not more accurate when using a compiled unitary that matches the isometry better than the optimization tolerance. For NISQ devices in particular, fine-tuning of isometry properties through the introduction of additional entangling gates may in fact produce worse results due to the increased noise in the circuit compared to a shallower representation. These properties have motivated us to pursue optimizations of the tensor network structure as well as a set of greedy compilation heuristics, inspired by Ref. [69], that we outline in what follows.

The key objects that we want to optimize in this section are the isometries $\hat{L}^{[i]}$ defined by the elements of the MPS in left-canonical form, see Eq. (4) of Sec. II A. Given that the binary encoding map used in this work, Eq. (1), is real-valued, all MPS tensors are real-valued, and this extends to the isometries. We will display the isometries using plots of the elements of their matrix representations in a fixed basis, as in Fig. 2. In this and similar plots, the basis ordering is defined with the physical qubit (i.e., the qubit that begins in the $|0\rangle$ state and is read out after each isometric operation) as the least significant qubit such that an isometry acting on a χ -dimensional ancilla $\alpha \in \{0, \dots, \chi - 1\}$ and a physical qubit $q \in \{0, 1\}$ has state indices

$$\text{index}(|\alpha q\rangle) = 2\alpha + q. \quad (12)$$

For isometries that have their ancilla states decomposed into qubits, we order those qubits $a_i \in \{0, 1\}$ such that significance increases with label index i , i.e.,

$$\text{index}(|a_{n_{\text{anc}}} \dots a_1 q\rangle) = \sum_{i=1}^{n_{\text{anc}}} 2^i a_i + q, \quad (13)$$

FIG. 2. *Example isometry for optimization.* An example isometry acting on a single physical qubit in the state $|0\rangle$ and a $\chi = 7$ -level ancilla, taken from the MNIST example in Sec. V. The real-valued matrix representation of this isometry is plotted in the basis with the physical qubit as the least significant qubit, see Eq. (12), and values not defined by the isometry are set to zero (white). The isometry has been cleaned to remove small numerical values resulting from classical optimization, but no further optimization has been applied.

The isometry in Fig. 2 acts on a physical qubit and a $\chi = 7$ dimensional ancilla, transforming the state $|00\rangle$ into a superposition of $|11\rangle$ and $|60\rangle$, the state $|10\rangle$ into $|00\rangle$, and so on. We note that the isometry in Fig. 2 is undefined when acting on states with $|q = 1\rangle$ in accordance with the sequential preparation scheme, but takes arbitrary ancilla states as inputs. Because of the isometry property, we only need to account for the nonzero elements of the operation when matching to a unitary, and so do not need to distinguish between zero elements and undefined elements.

As a first step in compilation, we will want to “clean” the isometries from the classical model in order to remove noise at the level of the classical optimization tolerance, otherwise we will expend effort attempting to compile this noise into quantum operations that will not improve the fidelity of the calculation. This amounts to implementing a filter on the MPS to remove elements below some tolerance level ε , which can be accomplished by using MPS compression to find the MPS with specified resources (e.g., restricted bond dimension χ) $|\phi\rangle$ that is closest in the L_2 -norm to a target MPS $|\psi\rangle$ that has higher resource requirements (χ'). While this is optimally done variationally [19], a simple and practical method for performing this operation is to use local SVD compression, in which the MPS tensor of the orthogonality center $A^{[i]}$ is decomposed by the SVD as

$$A_{\alpha\beta}^{[i]j_i} \rightarrow \sum_{\mu} U_{(\alpha j_i)\mu} S_{\mu} V_{\mu\beta}, \quad (14)$$

$$A_{\alpha\beta}^{[i]j_i} \rightarrow \sum_{\mu} U_{\alpha\mu} S_{\mu} V_{\mu(j_i\beta)}, \quad (15)$$

where the upper expression is for a right-moving update and the lower for a left-moving update. We can truncate the bond dimension by keeping only the χ largest singular values, or determine the new bond dimension implicitly through a

singular value cutoff ε as

$$1 - \sum_{\mu=1}^{\chi} S_{\mu}^2 / \sum_{\mu} S_{\mu}^2 < \varepsilon. \quad (16)$$

When the MPS tensor is the orthogonality center, this condition is equivalent to a L_2 -norm optimization of the full wave function. Replacing $A^{[i]}$ by the truncated U for a right-moving update or by V for a left-moving update and contracting the truncated SV or US into the neighboring tensor completes the local optimization. Sweeping the optimization across all tensors completes the filtering step. Since the optimization only deals with the parameters of a single MPS tensor at a time, it is not guaranteed to be globally optimal, but this simple procedure works well in practice. As a side benefit, ending the optimization by applying the update Eq. (14) and replacing the MPS tensor $A^{[i]}$ with U for each tensor places the MPS in left-canonical form, from which the isometries for sequential preparation can be constructed from the tensor elements [see Eq. (2)].

A. Ancilla permutation and the diagonal gauge

The conversion of an MPS into left canonical form uses the gauge freedom inherent in MPSs to ensure that each of the MPS tensors $\mathbb{L}^{[i]j_i}$ satisfies the isometry constraint

$$\sum_{\alpha j} L_{\alpha\beta}^{[i]j} L_{\alpha\beta'}^{[i]j*} = \delta_{\beta\beta'}, \quad (17)$$

without changing the overall quantum state. This gauge freedom specifies that any invertible matrix $\mathbb{X}$ and its inverse can be placed between any two tensors of the MPS, i.e.,

$$\tilde{\mathbb{A}}^{[i]j_i} = \mathbb{A}^{[i]j_i} \mathbb{X}, \quad (18)$$

$$\tilde{\mathbb{A}}^{[i+1]j_{i+1}} = \mathbb{X}^{-1} \mathbb{A}^{[i+1]j_{i+1}}. \quad (19)$$

However, the constraint Eq. (17) still allows for the insertion of any *unitary* matrix and its inverse on either the left or right bond basis of an MPS tensor in left-canonical form $\mathbb{L}^{[i]j}$ without changing the state or the isometry conditions. This freedom stems from the fact that the bond degrees of freedom are only used to mediate correlations between the physical degrees of freedom and are not directly measured, and so have no preferred basis for representation. We can attempt to exploit this freedom to produce MPS models that are more amenable to compilation on a given target hardware. We note that, just as with the ordinary gauge freedom of MPSs, a change of gauge affects two neighboring MPS tensors at a time, and so an operation that may benefit one tensor also affects its neighbors and so on down the network. Thus, the optimal choice of gauge requires a global optimization across all tensors.

To utilize the ambiguity in the basis representation of the ancilla states, we have devised a simple procedure that we have found to aid in compiling isometries for QAML models. The heuristic guiding our scheme is to ensure that operations are as “diagonal” as possible, in the sense that qubits preferentially remain in their same state rather than being swapped or mixed with other ancilla qubits. Operationally, in order to work only within the ancilla basis where we have freedom of

representation, we define a matrix of overlaps

$$M_{\alpha\beta}^{[i]} = \sum_j L_{\alpha\beta}^{[i]j*} L_{\alpha\beta}^{[i]j}, \quad (20)$$

which “integrates out” the physical qubit from the isometry used for sequential preparation, and so acts only in the ancilla space. A diagonal $\mathbb{M}^{[i]}$ is desired, as this would perfectly preserve the individual ancilla basis states and so reduce the number of quantum operations required. Recalling that we are only changing either the left or right basis of $\mathbb{M}^{[i]}$ at a time, one possible option to increase its diagonal dominance through transformation of either the left or right basis is to use the polar decomposition $\mathbb{M}^{[i]} \rightarrow \mathbb{U}^{[i]} \mathbb{P}^{[i]}$ or $\mathbb{M}^{[i]} \rightarrow \mathbb{P}^{[i]} \mathbb{U}^{[i]}$ with $\mathbb{U}^{[i]}$ unitary and $\mathbb{P}^{[i]}$ Hermitian and positive semidefinite. Using $(\mathbb{U}^{[i]})^{1/2}$ to transform the bases of $\{\mathbb{L}^{[i]j_i}, j_i = 1, \dots, d\}$ would transform $\mathbb{M}^{[i]}$ into $\mathbb{P}^{[i]}$; however, this transformation does not preserve sparsity in the $\mathbb{L}^{[i]j_i}$, and we have found that it often leads to more complex operators in practice. Instead, we use the values of $\mathbb{U}^{[i]}$ from the polar decomposition to define a permutation of the ancilla basis states as, e.g.,

$$\tilde{L}_{\alpha, \arg\max_{\beta} |\mathbb{U}_{\beta}^{[i]}|}^{[i]j} = L_{\alpha\beta}^{[i]j}. \quad (21)$$

This operation does preserve sparsity, and results in more diagonal operations in the ancilla degrees of freedom. An example of the isometries for a QAML model with and without this permutation procedure are shown in the right and left panels of Fig. 3, respectively. We see that the permutation of the basis states does result in a more diagonal isometry operator, as desired. A procedure for fixing the overall sign of the isometries, as well as additional ambiguities that can arise for tensors near the end of the representation, are discussed in Appendix B.

As with the usual transformation of MPS gauge to mixed canonical form [19], there is a “right-moving” update that permutes the right bond basis of a tensor $\mathbb{A}^{[i]}$ and the left bond basis of $\mathbb{A}^{[i+1]}$ and a “left-moving” update that permutes the left bond basis of $\mathbb{A}^{[i]}$ and the right bond basis of $\mathbb{A}^{[i-1]}$. When applied to all tensors, we say that the MPS is in the *diagonal gauge*, as it is the gauge which enforces the isometries for state preparation to be as diagonal as possible (according to our particular cost functions). We stress that the MPS is still in left-canonical form, and so the sequential preparation scheme still holds; the diagonal gauge merely uses the unitary freedom remaining in the left-canonical form to further optimize the state preparation procedure while maintaining sparsity. There is a single tensor that is not optimized at a certain location k in the transformation to the diagonal gauge that we call the *diagonality center*, analogous to the orthogonality center of mixed canonical form. While the location of the diagonality center can again be used as an optimization parameter, we have found it convenient to set the diagonality center to an isometry that is initially an identity matrix. Such an isometry can always be introduced by padding the classical data vectors with a zero at location k . The reason for our choice is that the permutation to diagonal gauge will transform this identity isometry into a permutation matrix, which is likely to be easier to compile with high fidelity than a general, nonsparse

FIG. 3. Application of diagonal gauge. Example isometries for an operations with a $\chi = 7$ dimensional ancilla before (top: same isometry as Fig. 2) and after (bottom) applying the diagonal gauge transformation Eq. (21) to the right ancilla basis states.

isometry. Specific techniques for compilation will be presented in a later section.

B. Greedy compilation heuristics

Following the fixing of gauge outlined in the last subsection, we are in a position where we now want to transform the isometries $\hat{L}^{[i]}$ [see Eq. (4)] into operations to be performed on quantum hardware. The target hardware will have a collection of qubits laid out with a given topology and an allowed gate set of single-qubit rotations and entangling gates between pairs of qubits. Generally speaking, two-qubit gates are subject to higher degrees of noise than the single-qubit gates, and so higher-fidelity operations will be obtained by using as few two-qubit gates as possible. As an example, the error map and qubit/gate topology for the IBMQ-X2 machine is shown in Fig. 4. For this device, the single-qubit gates are defined by [4]

$$\hat{U}_3(\theta, \phi, \lambda) = \begin{pmatrix} \cos \frac{\theta}{2} & -e^{i\lambda} \sin \frac{\theta}{2} \\ e^{i\phi} \sin \frac{\theta}{2} & e^{i(\lambda+\phi)} \cos \frac{\theta}{2} \end{pmatrix}, \quad (22)$$

and the two-qubit gates are controlled-NOT (CNOT) gates, which are allowed only between qubits designated with a solid line in Fig. 4. As shown in the figure, the average error of the CNOT gates at the time of this measurement was $\sim 2.6\%$, while the error of the single-qubit gates was $\sim 0.15\%$. Hence, a goal in compiling our isometries is to use as few gates as

FIG. 4. Exemplar NISQ hardware architecture. The qubit layout (circles), CNOT coupling topology (lines), and error sources of the five-qubit IBMQ-X2 device [4] as an exemplar NISQ machine. We note that these error rates are a snapshot, and are subject to fluctuations.

possible, and especially to minimize the number of two-qubit gates.

In our compilation heuristic, we enumerate possible unitaries by constructing a tree of potential circuit structures with continuous parameters to be optimized. The root node of our tree is comprised of a single-qubit gate [such as the $\hat{U}_3$ gate in Eq. (22)] for each qubit. Each node in the tree has a child node corresponding to the placement of an entangling gate in one of its allowed positions, and then adding single-qubit gates to the qubits acted on by the entangling gate. Any circuit that can be constructed using the allowed entangling gates and single-qubit rotations corresponds to a node in this tree, as proved in Ref. [69]. In order to select between nodes in this tree, we define a cost function between the isometry $\hat{L}$ and a unitary candidate $\hat{U}$

$$\mathcal{C}(\hat{U}, \hat{L}) = \sum_{(i,j) \in \mathcal{S}} |U_{i,j} - L_{i,j}|^2, \quad (23)$$

in which $\mathcal{S}$ denotes the set of indices such that the elements of the matrix representation of the isometry are greater than some tolerance $|L_{i,j}| > \delta$. Because of the isometry property of $\hat{L}$ and the unitarity of the candidate gates $\hat{U}$, we can optimize only over the elements in $\mathcal{S}$, which reduces the computational complexity of the cost function. Our optimization will select a particular unitary $\hat{U}$ as being acceptable when the cost function drops below a specified tolerance ε .

The optimization procedure begins by optimizing the root node (single-qubit gates) over its parameters and checking the cost function if an acceptable gate is found. If no acceptable gate is found, a queue of gates corresponding to adding an entangling gate and a pair of single-qubit gates to the root node in all allowed locations as outlined above is formed, and these gates are optimized and their cost functions recorded. If no gate from this queue is acceptable, a priority queue is formed by sorting the gates from this set according to their cost functions and then appending entangling gates and single-qubit rotations as above. In order to avoid an exponential growth

FIG. 5. Comparison of greedy gate compilation procedure with methods of Ref. [70]. (a) Target isometry, which can be completed to a permutation operator. (b) Matrix plot of result from greedy compilation procedure (cost function $\sim 2 \times 10^{-15}$). (c) Matrix plot of result from the methods of Ref. [70]. (d) Quantum circuit representation of greedy compilation procedure result. (e) Quantum circuit representation of Ref. [70] result.

of the number of search considerations, we limit the number of gates forming the starting point of the priority queue (i.e., before appending new entangling gate and single-qubit rotations) to a fixed number. This number is used as a convergence parameter, and can vary between optimization cycles; we find that it is useful to allow more gates in early optimization cycles where the operations involve fewer parameters and so optimization is fast, and then to decrease the number of kept gates as the circuits become deeper. Further details on our implementation of this procedure and some problem-specific implementations are provided in Appendix C.

Several “generic” methods for the compilation of isometries exist, as reviewed in, e.g., Ref. [71], which can be used as a baseline for comparison. These algorithms also underlie the implementation in QISKIT [4]. In the generic approach, the matrix representation of the isometry is decomposed, e.g., a single column at a time or by the cosine-sine decomposition, and the resulting decompositions expressed in terms of multiqubit controlled operations, which are themselves decomposed into a target gate set using known representations. These approaches are constructive, and so will find decompositions of any isometry in principle, but they are not designed to find the most efficient representation by some metric, e.g., the number of entangling gates. Further, as noted above, the use of such generic algorithms requires an “isometric completion” in the case that the bond dimension χ is not a power of 2, and may expend additional resources in exactly compiling noise in the isometries. Special purpose methods have also been developed for compiling permutation gates in Ref. [70], which have been shown to outperform the generic algorithms in some cases. This method uses a reversible logic synthesis

to map the permutation into a reversible circuit comprised of single-target gates, and these single-target gates are then compiled into networks of CNOTs, Hadamard gates, and $\hat{R}_z(\theta) = |0\rangle\langle 0| + e^{i\theta}|1\rangle\langle 1|$ rotations.

In order to compare our methods with the generic, constructive method for compiling isometries, we consider the isometry shown in Fig. 3. As noted above, in order to utilize the generic methods we have to map this isometry into a complete isometry over a set of qubits, which requires us to define the action of the isometry on the state in which the ancilla qubits are all in the state $|1\rangle$, which was left unconstrained by the optimization procedure. For simplicity, we use the “isometric completion” in which the operator takes this state to itself without modifying the state of the physical qubit. Using the `iso` method of the QuantumCircuit class from Qiskit [4] implementing the generic methods of Ref. [71] on the unconstrained `ibmq.qasm_simulator` hardware topology produces a gate representation with 122 CNOTs at `optimization_level` 0, and 120 CNOTs at `optimization_level` 3. The greedy compilation procedure presented in this work achieves a representation with a cost function error of 5.6×10^{-10} with an order of magnitude fewer entangling gates for this particular isometry. An explicit circuit representation is given in Fig. 36(d) of Appendix E.

As a point of comparison for the specialized methods for permutation gates studied in Ref. [70], we consider the isometry shown in Fig. 5(a). This is indeed a permutation on the space acted upon, and so can be represented by a family of “unitary completions.” We take the straightforward choice of unitary completion in which we leave the ancilla qubits unchanged by the permutation, as shown in Fig. 5(c).

The result of applying our greedy compilation procedure is given in matrix form in panel (b), and in quantum circuit form in panel (d). This gate requires seven CNOTs, and has a cost function error of $\sim 2 \times 10^{-15}$. The result of applying the methods of Ref. [70] are shown in panels (c) and (e); the gate here requires 14 CNOT operators. Generally speaking, we find that our greedy compilation procedure finds comparable or better gates for isometries corresponding to near-diagonal permutations compared to using the methods of Ref. [70] with the straightforward unitary completion given above. However, it is also worth noting that our procedure is designed for isometries and so generally does not produce permutation operators on the entire space at the end of optimization. That is to say, the optimized gate is a permutation in the space spanned by the isometry, but the full unitary is not a permutation, see, e.g., Fig. 35 of Appendix E. It is also worth noting that for complex, highly nondiagonal permutations, as can occur for the diagonality center when transforming to the diagonal gauge, the methods of Ref. [70] can produce more efficient representations.

IV. EXACTLY SOLVABLE BENCHMARK MODEL

As an exactly solvable benchmark, we consider an MPS Born machine encoding the probability distribution of classical discrete data vectors $\mathbf{x}$, $x_i \in \{0, 1\} \forall i$. The simplest nontrivial situation is when the data vectors consist of all zeros except for a single 1. Let us denote the probability that the 1 resides at location i as p_i , with $\sum_{i=0}^{N-1} p_i = 1$. It can be shown that this data can be represented exactly as a bond dimension 2 MPS Born machine with tensors

$$A_{00}^{[0]0} = 1, \quad A_{01}^{[0]1} = e^{i\phi_0} \sqrt{p_0}, \quad (24)$$

$$A_{00}^{[j]0} = 1, \quad A_{01}^{[j]1} = e^{i\phi_j} \sqrt{p_j}, \quad A_{11}^{[j]0} = 1, \quad (25)$$

$$A_{10}^{[N-1]0} = 1, \quad A_{00}^{[N-1]1} = e^{i\phi_{N-1}} \sqrt{p_{N-1}}, \quad (26)$$

with the $\{\phi_j\}$ denoting arbitrary phases. The presence of a large number of arbitrary phases is a generic feature of TN models for generative applications: since the square of the wave function is used to generate classical data samples, the phase structure of the wave function is generally underconstrained. This in turn implies that TN models can have some flexibility over the particular gate set used to entangle the physical qubits to the ancillae without affecting the sampling outcomes. The exactly solvable model encapsulated by Eqs. (24)–(26) is a useful benchmark both because it is the simplest nontrivial example of a sequentially preparable QAML model, involving a single ancilla qubit, and because it can be exactly solved for any classical data vector length and probabilities $\mathbf{p}$. An example dataset for $\mathbf{p} = (1/5, 1/20, 1/20, 1/4, 1/5, 1/4)$ is given in Fig. 6.

The construction in Eqs. (24)–(26) is reminiscent of the well-known MPS representation of the W state [35]. In order to convert this generic MPS into a sequential qubit preparation scheme we should place the MPS into left-canonical form. Since the bond dimension is known, we can do so in terms of the QR decomposition. The explicit calculation is relegated to Appendix D, where we also discuss a “hand compiled” ver-

FIG. 6. Example dataset for an exactly solvable MPS generative model with a six-dimensional probability vector $\mathbf{p} = (1/5, 1/20, 1/20, 1/4, 1/5, 1/4)$.

sion of the resulting isometries to be used as a benchmark for our automated procedures. The outcomes of the calculation are the exact isometries

$$\hat{L}^{[N-1]} = (\cos \theta_{N-1} |1_a 0_q\rangle + e^{i\phi_{N-1}} \sin \theta_{N-1} |0_a 1_q\rangle) \langle 0_a 0_q|, \quad (27)$$

$$\hat{L}^{[j]} = |0_a 0_q\rangle \langle 0_a 0_q| + (e^{i\phi_j} \sin \theta_j |0_a 1_q\rangle + \cos \theta_j |1_a 0_q\rangle) \langle 1_a 0_q|, \quad (28)$$

in which $\cos \theta_j = \sqrt{\frac{\sum_{i \leq j} p_i}{\sum_{i \leq j} p_i}}$, $\sin \theta_j = \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}$, and $|x_q\rangle$ and $|y_a\rangle$ denote the states of the physical (sampled) qubit and the ancilla qubit, respectively. A natural unitary completion of the operator $\hat{L}^{[j]}$ is given in the basis $\{|0_a 0_q\rangle, |0_a 1_q\rangle, |1_a 0_q\rangle, |1_a 1_q\rangle\}$ as

$$[\hat{U}^{[j]}] = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta_j & e^{i\phi_j} \sin \theta_j & 0 \\ 0 & -e^{-i\phi_j} \sin \theta_j & \cos \theta_j & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (29)$$

which can be compiled exactly using the circuit shown in Fig. 7.

A. Training and compilation

In this section, we detail the application of the methods outlined in this paper to the exactly solvable benchmark in Sec. IV using the probabilities $\mathbf{p} = (8/31, 18/31, 5/31)$. To demonstrate our methods to this benchmark case, we train a $\chi = 2$ Born machine using the single-site gradient descent described in Sec. II B and compile it into gates using the procedures of Sec. III with the diagonality center at 1 and a greedy optimization tolerance of 5×10^{-4} . Following the usual parlance of MPSs from condensed matter physics, we will refer to the physical indices of the MPS tensors as sites. The results of this procedure are shown in Fig. 8, with panels (a)–(c) for site 0, panels (d)–(f) for site 1, and panels (g)–(i) for site 2. The final cost functions for sites 0, 1, and 2 are 6.7×10^{-9} , 7.3×10^{-10} , and 2.0×10^{-9} , respectively. As a point of comparison, we consider a “hand compiled” version of the unitary completed isometries Eqs. (29). Taking $\phi = 0$, we can compile these gates using the circuit shown in Fig. 7.

FIG. 7. Circuit decomposition for the natural unitary completion $\hat{U}^{[j]}$ using a gateset of single qubit rotations and CNOTs. The upper line is the physical qubit and the lower line is the ancilla.

FIG. 8. Exactly solvable Born Machine benchmark isometries. Isometries and optimized gates for the three-site exactly solvable Born Machine benchmark. (a)–(c) are for site 0, (d)–(f) are for site 1, and (g)–(i) are for site 2. (a), (d), and (g) are the isometries output from the classically trained model, (b), (e), and (h) are matrix plots of the unitaries output by our greedy compilation procedure, and (c), (f), and (i) are circuit representations of the optimized unitaries.

However, we additionally note that with the assumption that the physical qubit starts in the state $|0_q\rangle$, the first CNOT in Fig. 7 is the identity, and so can be neglected, leading to a circuit with three CNOTs. We see that the obtained quantum circuits are substantially different than those obtained by hand compilation of the “natural” unitary completion, but are still of very high fidelity in the space spanned by the isometry. In addition, the gates for sites 0 and 1 are shallower than the hand-compiled gate, which may be anticipated based on known optimality results for two-qubit gates [72].

B. Performance of benchmark on cloud-based hardware

In this section, we present results for the exactly solvable benchmark model running on cloud-based NISQ hardware, using IBM devices as an example. We note that the current IBM hardware does not allow measurement and re-initialization during an experimental run, and so our sequential preparation schemes cannot be directly implemented on these devices. However, we can still test our generative models by implementing the gates $\hat{U}^{[i]}$ of the sequential preparation scheme on a register of $(N + 1)$ qubits prepared in the $|0 \dots 0\rangle$ state, coupling each physical qubit to the same ancilla in order from $(N - 1)$ down to 0. This procedure is limited in practice by the number of available qubits and their connectivity to a single ancilla qubit. However, for devices with a cross-shaped topology, such as the IBMQ-X2, we can readily couple up to 4 qubits to a central ancilla qubit, and for devices with a T-shaped topology, such as the Vigo, we can couple up to 3 qubits to a single ancilla.

Utilizing this approach for the exactly solvable Born Machine model with the probability vector given in Sec. IV A, we find the circuits shown in Fig. 9. Here, the physical qubits (those that are sampled to obtain output classical data vectors) are assigned to be qubits 0, 1, and 3, and the ancilla is qubit 2. The upper panel is for the hand-compiled circuits from Fig. 7, and the lower panel is the circuit from Fig. 8 using the workflow put forth in this work. The dashed vertical lines demarcate the circuits corresponding to the individual sites of the Born machine, but are inessential and neighboring single-qubit rotations can be joined for increased efficiency.

As metrics for assessing the performance of our QAML models, we utilize both the raw experimental counts used to infer measurement probability distributions and a convex version of the Kullback-Leibler (KL) divergence between the ideal (p_T) and estimated (p_N) distributions, as implemented in SCIPY [73],

$$\text{KL}(p_T, p_N) = \begin{cases} p_T [\ln(\frac{p_T}{p_N}) - 1] + p_N & p_T > 0, p_N > 0 \\ p_N & p_T = 0, p_N \geq 0 \\ \infty & \text{otherwise} \end{cases} \quad (30)$$

The noise levels of NISQ devices fluctuate over time, and so to account for these statistical variations we implemented a jackknife procedure [74] for the mean and variance including bias correction, utilizing 25 experimental runs per day of $2^{13} = 8192$ shots each across 5 days. We further refine each experimental run using the measurement noise filter implemented in QISKIT [4], which produces a measurement noise correction map from a collection of calibration measurements which are performed immediately before the experimental shots.

The results of our jackknife analysis on the experimental measurement counts per state are shown in Fig. 10. Here, panels (a) and (c) are the results for the hand-compiled model circuit in Fig. 9(a) and panels (b) and (d) are for the auto-compiled circuit in Fig. 9(b). Panels (a) and (b) are run on the IBMQ-X2 device, and panels (c) and (d) are on the IBMQ-Vigo device. In all panels, the rightmost green bar represents the ideal counts given by the model probability vector $\mathbf{p} =$

FIG. 9. Comparison of hand-compiled and auto-compiled circuits for exactly solvable test case. The exactly solvable benchmark with three physical qubits (0, 1, and 3) and a single ancilla qubit (2) implemented as a quantum circuit using the hand-compiled circuits in Fig. 7 (top) or the autocompiled gates in Fig. 8 (bottom).

(8/31, 18/31, 5/31), the center blue bars are the raw experimental measurements without noise calibration applied, and the leftmost orange bars are the experimental measurements with the noise calibration applied. The black lines centered on

the tops of the bars indicate the 1σ confidence intervals from the jackknife procedure. As noted above, qubits 0, 1, and 3 map to the probabilities p_0 , p_1 , and p_2 , respectively, and qubit 2 is the ancilla. Clearly, the application of the measurement

FIG. 10. Comparison of ideal, measured, and noise-corrected measurement outcomes on quantum hardware. The results of the hand-compiled benchmark model shown in Fig. 9(a) are jackknifed over several days of independent experimental runs on the IBMQ-X2 [panel (a)] and IBMQ-Vigo (c) hardware. Similarly, the jackknifed results for auto-compiled benchmark model shown in Fig. 9(b) are shown for the IBMQ-X2 (b) and IBMQ-Vigo (d) hardware. The rightmost green bars are the noiseless expectations, the center blue bars are the raw measurements, and the left orange bars have a measurement filter applied. Black lines indicate 1σ confidence intervals.

FIG. 11. Convex KL divergence between ideal and measured QAML outcomes with time. The convex KL divergence Eq. (30) between the ideal, noiseless measurement probabilities and the measurement probabilities inferred from 25 experiments of 2^{13} shots is shown as a function of experimental run day for the IBMQ-X2 (top) and IBMQ-Vigo (bottom) devices. Filled symbols use the raw experimental counts (center blue bars in Fig. 10) and empty symbols use the counts with measurement noise filter applied (left orange bars in Fig. 10). Lines indicate the KL divergences computed using all measurements from all days.

noise filter improves the fidelity of the results. Also, generally speaking, the results for the Vigo device (lower panels) are closer to the ideal results than for the X2 (upper panels). The largest probability state resulting from errors is the state $|0000\rangle$ with no “hot” physical bits, followed by $|1100\rangle$, with the two highest probability physical bits “hot.” We note that the outcomes involving the ancilla qubit in the $|1\rangle$ state can be removed in postselection by virtue of the fact that the sequential preparation scheme should end with the ancilla in the $|0\rangle$ state (see Sec. II A), but this results in small corrections for the present case. Finally, we see that the autocompiled results using the approach of Sec. III (right panels) are generally closer to the ideal results than the hand-compiled circuits (left panels), though this is not true for each state individually.

In Fig. 11, we display the KL divergence between the ideal, noiseless probabilities of measuring each individual quantum state and the measurement probabilities estimated from 25 experiments of 2^{13} shots without (filled symbols) and with (empty symbols) the measurement noise calibration filter applied. The x axis denotes consecutive experimental days, and the horizontal lines indicate the KL divergence re-

FIG. 12. Example processed MNIST data produced by down-sampling through a max filter to 7×7 pixels and binarization. Clockwise from top left, the truth labels are 5,0,4,1,9,4,1,3,1,2.

sulting from the distributions averaged over all days. Clearly, the application of the measurement noise filter improves the estimation of probabilities, as indicated by a lower KL divergence with respect to the ideal results. In addition, the autocompiled circuits (squares) show a lower KL divergence than the hand-compiled circuits, likely due to their shallower circuits. Finally, we find that the Vigo results in panel (b) have lower KL divergence than the X2 results in panel (a), indicating an overall lower noise level for these days, in spite of the day-to-day fluctuations in the KL divergence being comparable in magnitude between the two machines.

V. EXAMPLE USING THE MNIST DATASET

The exactly solvable benchmark presented in Sec. IV provided data that was simple and well-structured enough that it could be exactly memorized using a single qubit to mediate bistring correlations. In this section, we consider a QAML benchmark that is again a generative MPS Born machine, analogous to Sec. IV, but using data from the MNIST handwritten digit dataset [25], a canonical ML test case. The MNIST dataset consists of greyscale images, each consisting of 28×28 pixels, of the numbers 0 through 9. We process the data by passing through a filter that returns the max value from every contiguous 2×2 pixel block twice, resulting in images of size 7×7 . While this is not necessary, and produces less raw data available for learning, it reduces the number of isometries in the sequential preparation scheme to compile, allowing for both a more detailed case-by-case analysis and reducing the overall gate depth for the ancilla in the sequential preparation scheme. Our next step in processing this data is to binarize the greyscale images, such that we can use the binary qubit encoding for simplicity. Examples of the processed data are shown in Fig. 12.

In this work, we explore this dataset in the small-data regime where MPS models of modest bond dimension can memorize all patterns, analogous to the benchmark in Sec. IV, but the bitstrings demonstrate more complex correlations that require more quantum resources. The bitstrings are specified by mapping the 7×7 pixel arrays into binary vectors of length 49. As an example of the classical optimization procedure, an MPS Born machine with $\chi = 8$ (three ancilla qubits) converges to a negative log-likelihood of 2.563 following roughly 400 iterations of single-site gradient descent with a learning rate of $\eta = 10^{-4}$ (see Sec. II B) on the $N_T = 10$ item dataset shown in Fig. 12. Using $\chi = 16$ (four ancilla qubits), we reach the theoretical minimum value of the negative log-likelihood

FIG. 13. Convergence of classical TN optimization. The convergence behavior of the negative log-likelihood is shown for $\chi = 8$ (red solid lines) and $\chi = 16$ (blue dashed lines) in (a). Both results use single-site gradient descent with a learning rate of $\eta = 10^{-4}$. (b) displays sample data drawn from the initial, random $\chi = 16$ MPS model before optimization, (c) shows samples drawn from the $\chi = 8$ MPS model after optimization, and (d) shows samples drawn from the $\chi = 16$ MPS model after optimization. Because of the small size of the dataset, the $\chi = 16$ model is able to reach the theoretical minimum and so memorize the full dataset.

of $\ln N_T$ after roughly 200 iterations. The convergence behavior of these log-likelihoods is shown in Fig. 13, together with samples drawn from the model before and after optimization. Because the $\chi = 8$ model does not reach the theoretical minimum, data elements not seen in the training data are present in the samples. In contrast, the $\chi = 16$ model reaches the theoretical minimum, and so only produces data samples from the training data.

Following the classical optimization of the MPS tensors, we clean the MPSs to remove small numerical values from the classical optimization procedure, place it into left-canonical form to describe a sequential preparation scheme, and then fix the remaining permutation ambiguity in the bond degrees of freedom using the transformation to diagonal gauge described in Sec. III A. For the case in which the diagonality center is at site 35, we find the isometries collected in Figs. 20–68 of Appendix E. We compile these isometries using our greedy compilation procedure with a cost function [Eq. (23)] tolerance of $\varepsilon = 5 \times 10^{-4}$ except for the diagonality center, which we compile using the methods of Ref. [70] following the straightforward unitary completion procedure defined in Sec. III B; the results are again shown in Figs. 20–68 of Appendix E. In these figures, the “raw” circuits may include single-qubit rotations with rotation angles near zero and non-native parameterized gates as described in Appendix C, while the “expanded and cleaned” circuits compile the non-native

FIG. 14. Example data sampled from MNIST model on simulated noiseless quantum hardware. The data samples with probability $\geq 1\%$ obtained from the sequential preparation procedure defined by the isometries in Figs. 20–68 are displayed together with their probability of occurrence estimated from 2^{13} shots.

gates into single-qubit gates and CNOTs, collect adjacent single-qubit rotations, and then remove small single-qubit rotations with additional optimization passes. We note that $\varepsilon = 5 \times 10^{-4}$ translates into roughly $\sqrt{\varepsilon} \sim 2\%$ error in the elements of the compiled unitary, which is comparable to the entangling gate error rates of current cloud-based machines.

To investigate the fidelity of the compiled model, we implemented the sequential preparation procedure in which a single data qubit is coupled to three ancilla qubits using the isometries in Figs. 20–68 on the IBM qasm hardware simulator. As described in Sec. II A, the isometries are applied from site 48 down to site 0 with data qubit measurement in the z basis and reinitialization in the $|0\rangle$ state between the application of isometries. The outcomes of these measurements constitute a data sample of the model. While measurement and reinitialization in the midst of an experimental run are not supported on the current IBM quantum hardware, this operation is supported in the hardware simulators. Example data generated from 2^{13} runs on ideal, noiseless hardware is shown in Fig. 14. Here, the data samples with probability $\geq 1\%$ are shown together with their probabilities. The training data from Fig. 12 are clearly recognized as the elements with highest probability, save for the digit 5, whose highest probability data sample involves confusion with the 1 digit. We recall that deviations from the ideal training digits and their ideal occurrence probabilities of 10% are a result of the restriction of the model to $\chi = 8$ as well as the finite optimization tolerance ε . Comparing with Fig. 13(c), which shows samples taken from the classically trained model, it appears that the restriction to $\chi = 8$ has a greater influence on the probabilities and sample variations than the finite compilation tolerance.

To investigate the effects of hardware noise, we use a simple model of depolarization noise in which we take $\xi = 1 - F$ to be the average gate error, with F the average gate fidelity.

FIG. 15. Convex KL divergence between sampled data without and with hardware noise. The convex KL divergence Eq. (30) between a noiseless and noisy simulation is shown as functions of the CNOT error rate ξ_2 and the readout error ζ . The probability distributions of both the noiseless and noisy models are estimated from 2^{13} shots using a hardware simulator. The vertical dashed line indicates 1% CNOT error rate, a rough measure of the current state-of-the-art for NISQ devices.

With this, the depolarizing channel is represented by the operator

$$\hat{\mathcal{E}}_{\text{dep}} = (1 - p)\hat{I} + p\hat{D}, \quad (31)$$

in which $p = 2^{N_q}\xi/(2^{N_q} - 1)$ with N_q the number of qubits and the Kraus representation of the depolarizing channel is given by the operators

$$\hat{\mathcal{E}}_{\text{dep}} = \{\sqrt{1 - (4^{N_q} - 1)p/4^{N_q}}\hat{I}^{\otimes N_q}, \sqrt{p/4^{N_q}}\hat{P}\} \quad (32)$$

in which

$$\hat{P} = \{\hat{I}, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}^{\otimes N_q} \setminus \hat{I}^{\otimes N_q}, \quad (33)$$

is the set of N_q -qubit products of Pauli operators without the N_q -qubit identity matrix. In our model, we assign the same error ξ_2 to all CNOT gates and an error $\xi_1 = \xi_2 \times 10^{-2}$ to all single-qubit gates in accordance with typical IBM hardware characteristics. In addition to the depolarizing error, we also include an uncorrelated-qubit (tensor product) readout noise model parameterized by [75,76]

$$P(0|0) = 1 - \zeta, P(1|0) = \zeta, \quad (34)$$

$$P(0|1) = 2\zeta, P(1|1) = 1 - 2\zeta, \quad (35)$$

in which $P(a|b)$ denotes the probability of obtaining measurement outcome a from a preparation of the quantum state $|b\rangle$, assumed identical across all qubits for simplicity.

We again characterize the difference in predictions between the ideal and noisy outcomes using the convex KL divergence defined in Eq. (30). As before, we estimate the probabilities from ensembles of $N_s = 2^{13}$ shots calculated using a hardware simulator for both the noiseless “truth” model and the noisy models. This divergence is shown in Fig. 15 as a function of the error parametrizations ξ_2 and ζ . The CNOT error rate ξ_2 parameterizes both the two-qubit and single-qubit depolarization errors and ζ parameterizes the readout error. The divergence shows a rapid rise driven by the appearance of data samples not present in the noiseless model. The samples

FIG. 16. Example data sampled from MNIST model on simulated noisy quantum hardware. The highest probability data samples from the sequential preparation procedure defined by the isometries in Figs. 20–68 on a simulated machine with depolarization error rate $\xi_2 = 0.01$ (dashed line in Fig. 15) are displayed together with their probability of occurrence estimated from 2^{13} shots.

with greatest occurrence drawn from the model evaluated at the CNOT error rate $\xi_2 \sim 0.01$, indicated by the dashed vertical line in Fig. 15, and $\zeta = 0$ are shown in Fig. 16. We can recognize many of the digits from the training set in this model, but they occur with significantly lower probabilities due to the appearance of additional noise-driven patterns. Because of the sequential preparation, we can expect that the bits near the end of the bitstring (i.e., for sites near 48) are produced at higher fidelity than those of lower site indices because of errors present in manipulating the ancilla qubits. We see that this is the case in Fig. 17, which displays the convex KL divergence as functions of the CNOT error rate

FIG. 17. Convex KL divergence variation with noise and bitstring length. The convex KL divergence Eq. (30) between a noiseless and noisy simulation is shown as a function of the CNOT error rate ξ_2 and the sampled bitstring length at zero readout error ($\zeta = 0$). The probability distributions of both the noiseless and noisy models are estimated from 2^{13} shots using a hardware simulator.

ξ_2 and the number of bits sampled in the bitstring at zero measurement error $\zeta = 0$. Nonmonotonic behavior is due to the significant differences in complexity of the gate sequences to produce individual bits, see Figs. 20–68 of Appendix E. A rise in error is seen as the number of bits increases as the total gate depth increases, peaking around ten bits. After this point the KL divergence levels off as additional bits may require shallower gate sequences, bringing the overall agreement between the noisy and true probability distribution closer.

VI. CONCLUSIONS AND OUTLOOK

We have presented a complete workflow for generative quantum-assisted machine learning (QAML) using Born machines with a matrix product state (MPS) tensor network (TN) structure. In our workflow, classical data is encoded into quantum states using an embedding map, the ensemble of quantum states is learned as a TN Born machine using a classical DMRG-like procedure with gradient descent of the negative log-likelihood, and the model is compiled into operations for target quantum hardware to obtain data samples as measurement outcomes. Using MPS-based models enables the use of highly quantum resource-efficient sequential preparation schemes requiring $\mathcal{O}(1)$ qubits for a classical data vector length N and $\mathcal{O}(\log_2 \chi)$ qubits for bond dimension χ , which encapsulates the model expressivity. This condition is also sufficient; any model that uses a sequential preparation procedure to generate samples using only $\mathcal{O}(1)$ qubits for a classical data vector length N can be expressed as an MPS and so our methods are applicable. We expect this class of maximally quantum resource-efficient models will be of paramount importance when benchmarking QAML applications on near-term devices.

We presented several optimizations in the compilation stage of our QAML workflow, such as the introduction of the diagonal gauge of the MPS model that utilizes inherent freedom in the model representation to reduce the complexity of the compiled model, as well as greedy heuristics for finding shallow gate sequences matching a target isometry to a specified tolerance given hardware topology and allowed gate constraints. We presented an exactly solvable benchmark model requiring two qubits, and assessed its performance on currently available quantum hardware. We also presented an example application modeling features extracted from the MNIST dataset parametrically with depolarizing and readout hardware noise using a hardware simulator.

Our results lay the groundwork for utilizing TN models in practical QAML applications, and leave several avenues for future research. First, the QAML demonstrations given in this work consist of overfit models, and so do not constitute “true” machine learning models which should be able to appropriately generalize from data. This is a result of either using data with very simple structure, as in our exactly solvable model, or using a very small sample size of training data, as in our MNIST application. Small sample sizes were used in the present work to enable detailed analysis of model performance with limited quantum resources. In future work, the generalization power of TN-based QAML models on NISQ hardware will be explored moving towards

the large-data regime. We also note that other studies have indicated that TN models with current training strategies generally have a tendency towards overfitting [77]. Second, we have focused on the applications of MPSs to generative modeling, but other TN structures, such as tree tensor networks [39,78] may also be useful for QAML applications, as well as other tasks such as feature extraction and classification. The procedures outlined in this paper can be readily adapted to compiling the isometries appearing in models for other TNs and other applications. Finally, the procedure outlined in this paper wherein a model is trained classically before being compiled to a quantum device cannot by itself yield a quantum advantage, as it requires the model to be both classically and quantumly simulable. However, our procedures will be useful in designing and analyzing TN-inspired model structures for scaling towards the classically intractable regime, and can also serve as “preconditioners” where a model trained using optimal classical strategies is augmented with additional quantum resources and then trained directly on the quantum device or in a hybrid quantum/classical optimization loop, potentially avoiding local minima and speeding up optimization times.

ACKNOWLEDGMENTS

We would like to thank Dave Clader, Giuseppe D’Aguanno, and Colin Trout for useful discussions and would like to acknowledge funding from the Internal Research and Development program of the Johns Hopkins University Applied Physics Laboratory.

APPENDIX A: PROCEDURE FOR CONVERTING AN MPS MODEL INTO A SEQUENTIAL PREPARATION SCHEME

The purpose of this section is to review the procedure for converting between matrix product state models and a sequential preparation procedure with a χ -level ancilla. To begin, we first consider that we have a register of N qubits with states $|j_i\rangle$, $i = 0, \dots, N-1$, $j_i = 0, 1$ in which we want to encode data and a χ -level ancilla $|\alpha\rangle$, $\alpha = 0, \dots, \chi-1$ that can be used to entangle the qubits. Starting at the “right” end of the system, we can initialize the $(N-1)^{\text{st}}$ qubit using an operator $\hat{L}^{[N-1]}$ defined as

$$\hat{L}^{[N-1]} = \sum_{\alpha, j_{N-1}} L_{\alpha}^{[N-1]j_{N-1}} |j_{N-1}\alpha\rangle \langle 00|, \quad (\text{A1})$$

in which the coefficients $L^{[N-1]}$ satisfy the isometry condition

$$\sum_{\alpha, j_{N-1}} L_{\alpha}^{[N-1]j_{N-1}} L_{\alpha}^{[N-1]j_{N-1}} = 1. \quad (\text{A2})$$

Clearly, if we start our qubit and ancilla system in the state $|00\rangle$, this operation transforms it into the (entangled) state $\sum_{\alpha, j_{N-1}} L_{\alpha}^{[N-1]j_{N-1}} |\alpha j_{N-1}\rangle$, and the isometry condition ensures that this state is normalized. Moving to the next qubit, we now entangle it with the ancilla using the operator

$$\hat{L}^{[N-2]} = \sum_{\alpha, j_{N-2}, \beta} L_{\beta\alpha}^{[N-2]j_{N-2}} |j_{N-2}\beta\rangle \langle 0\alpha|, \quad (\text{A3})$$

in which the coefficients $L_{\beta\alpha}^{[N-2]j_{N-2}}$ expressed as a matrix $\mathbb{L}^{[N-2]j_{N-2}}$ are subject to the isometry condition

$$\sum_j \mathbb{L}^{[N-2]j^\dagger} \mathbb{L}^{[N-2]j} = \mathbb{I}_\chi, \quad (\text{A4})$$

with $\mathbb{I}_\chi$ the $\chi \times \chi$ identity matrix. This operation now puts the system in the state

$$\begin{aligned} & \hat{L}^{[N-2]} \hat{L}^{[N-1]} |0_{N-2} 0_{N-1} 0_{\text{ancilla}}\rangle \\ &= \sum_{j_{N-2}, j_{N-1}, \alpha} [\mathbb{L}^{[N-2]j_{N-2}} \mathbb{L}^{[N-1]j_{N-1}}]_\alpha |j_{N-2} j_{N-1} \alpha\rangle. \end{aligned} \quad (\text{A5})$$

We follow this same logic for all subsequent qubits, defining isometric operators that entangle them to the rest of the system using the ancilla, until we reach qubit 1, which is attached using the isometric operator

$$\hat{L}^{[0]} = \sum_{i_0, \beta} |i_0 0\rangle \langle 0 \beta|. \quad (\text{A6})$$

This operator puts the full system into the state

$$\begin{aligned} & \hat{L}^{[0]} \dots \hat{L}^{[N-1]} |0_0 \dots 0_{N-1} 0_{\text{ancilla}}\rangle \\ &= \sum_{j_0 \dots j_{N-1}} \mathbb{L}^{[0]j_0} \dots \mathbb{L}^{[N-1]j_{N-1}} |j_0 \dots j_{N-1} 0_{\text{ancilla}}\rangle. \end{aligned} \quad (\text{A7})$$

Hence, in the last step, the qubit states decouple from the ancilla. The qubit state takes the form of an MPS with the additional constraint that each of the MPS matrices $\mathbb{L}$ satisfies the left-orthogonal condition Eq. (A4). The above procedure can readily be read in reverse; given a general MPS QAML model with bond dimension χ ,

$$\sum_{j_0 \dots j_{N-1}} \mathbb{A}^{[0]j_0} \dots \mathbb{A}^{[N-1]j_{N-1}} |j_0 \dots j_{N-1}\rangle, \quad (\text{A8})$$

we can convert it into a sequential qubit preparation scheme with a χ -dimensional ancilla by putting the MPS in left-canonical form. This transformation to left-canonical form can be done without loss of generality using a well-known procedure involving an orthogonal decomposition, e.g., the singular value or QR decomposition [19]. Thus the tensors appearing in an MPS, which could result from a classical training optimization, can be formally (i.e., modulo compilation into native quantum operations for a given hardware architecture) translated into operations for deployment on quantum resources.

The above prescription assumed the presence of a register of N qubits, but due to the sequential nature of the preparation this is unnecessary, and a single “physical” qubit together with the χ -level ancilla suffices, provided we are not measuring any multiqubit properties of the state. As an example, we will consider drawing a sample from an MPS wave function generative model with the binary map Eq. (1). In this application, we first couple the qubit and ancilla as in Eq. (A1) starting from both in the fiducial state $|0\rangle$. We then measure the qubit in the computational basis, record its outcome as x_{N-1} , and then return it to the fiducial $|0\rangle$ state while leaving the ancilla unmeasured.

We note that the ability to re-initialize a single qubit independent of the others is not universally available in present-day hardware, but has been demonstrated in, e.g.,

trapped ion platforms [64]. We then re-entangle the ancilla and qubit using the operator $\hat{L}^{[N-2]}$ defined in Eq. (A3), measure the qubit and record the outcome as x_{N-2} , and again return the qubit to the $|0\rangle$ state. This procedure is repeated with the other operations $\hat{L}^{[j]}$ until a complete set of N measurements $\mathbf{x}$ is made, which constitutes a data sample. This procedure is denoted graphically in Fig. 1(d). Clearly, this only requires a single “physical” or “data” qubit (i.e., the one that is sampled) independent of the input data size N , and the construction of the χ -level ancilla requires only $\log_2 \chi$ qubits. We stress that the scheme above is formal in the sense that it produces isometries acting on quantum resources without reference to their actual physical representation or other hardware constraints such as limited coherence time, connectivity, gate sets, etc.. The translation of these formal isometries into operations to be dispatched on a given target hardware are detailed in Sec. III.

APPENDIX B: RESOLVING AMBIGUITIES IN THE TRANSFORMATION TO THE DIAGONAL GAUGE

In this section we provide additional detail on the resolution of two ambiguities that arise when converting MPS isometries into the diagonal gauge, as described in Sec. III A. The permutation operation Eq. (21) is ambiguous whenever the matrix $\mathbb{U}^{[i]}$ derived from the polar decomposition of $\mathbb{M}^{[i]}$ [see Eq. (20)] has multiple elements in a column with the same absolute value. Recalling that our sequential MPS preparation scheme requires that the ancilla start and end in the vacuum state, we see that this occurs for tensors near the extremal values of the representation when an ancilla qubit is first utilized or an ancilla qubit is decoupled from the remaining qubits. In such cases, we use the following alternate procedure to decide between permutations. First, we enumerate all basis permutations resulting from these ambiguities for a given tensor $\mathbb{L}^{[i]j_i}$ and construct their associated isometries $\hat{L}^{(\zeta)}$, in which ζ indexes permutations. To decide between these permutations, we again would like to make this operator as “diagonal” as possible, in the sense of minimizing the number of qubit operations being applied. We construct a simple cost function as follows: for each state indexed by the ancilla state α and the physical qubit q as above, we convert the state index into its binary representation $\mathbf{b}$, which effectively maps the ancilla state onto a collection of $\log_2 \chi$ qubits. As an example, the states of a four-dimensional ancilla and a single physical qubit give the representations

$$\text{index}(|0, 0\rangle) = 0 \rightarrow (0, 0, 0), \quad (\text{B1})$$

$$\text{index}(|0, 1\rangle) = 1 \rightarrow (0, 0, 1), \quad (\text{B2})$$

$$\text{index}(|1, 0\rangle) = 2 \rightarrow (0, 1, 0), \quad (\text{B3})$$

$$\text{index}(|1, 1\rangle) = 3 \rightarrow (0, 1, 1), \quad (\text{B4})$$

$$\vdots, \quad (\text{B5})$$

$$\text{index}(|3, 1\rangle) = 7 \rightarrow (1, 1, 1). \quad (\text{B6})$$

We now calculate a distance between two basis states (α, j) and (α', j') with respective binary representations $\mathbf{b}$ and $\mathbf{b}'$ as $\mathcal{D}[(\alpha, j), (\alpha', j')] = (\sum_\mu |b_\mu - b'_\mu|)^2$. The term in parentheses counts the number of individual qubit “flips” required to

convert one of the states into the other, and the square strongly penalizes multiqubit coordinated flips. We then use the cost function

$$C_\zeta = \text{Tr}(|\mathbb{L}^{(\zeta)}| \mathbb{D}), \quad (\text{B7})$$

in which $\mathbb{D}$ is the matrix with $\mathcal{D}[\bullet, \bullet]$ as elements and $|\mathbb{L}^{(\zeta)}|$ is the matrix of absolute values of $\mathbb{L}^{(\zeta)}$, to choose from between the $\mathbb{L}^{(\zeta)}$.

In addition to this permutation ambiguity, there is also a sign ambiguity on each of the bond states of the isometry. We again use diagonal dominance to fix this sign ambiguity in the following fashion: when our diagonal gauge update sweep is proceeding towards tensor to the right, we are updating the right bond basis of the current tensor. Hence, we identify the nonzero element closest to the diagonal in each column, and reverse the sign of this element if it is negative. In order that we do not change the overall state, we multiply the elements in each corresponding row of the tensor to the right in the MPS description. Similarly, in a left-moving sweep, we identify the nonzero elements closest to the diagonal in each row, flip signs if this element is negative, and flip signs in the corresponding columns of the tensor to the left in the MPS representation. Following transformation to diagonal gauge, we fix the signs of all elements of the diagonality center (chosen, as above, to be a permutation operator) to be positive. Any sign flips that occur in this transformation are propagated through neighboring MPS tensors until a row or column with elements of mixed sign is identified. This procedure discourages the appearance of permutation operators whose elements are both positive and negative.

APPENDIX C: IMPLEMENTATION AND OPTIMIZATION DETAILS OF GREEDY COMPILATION HEURISTICS

In this section, we briefly note details of our implementation of the greedy compilation heuristics of Sec. III B, along with some problem-specific optimizations. Our subroutine for the cost function takes as input a vector of parameters θ , constructs a matrix representation of the parameterized gate sequence

$$\hat{U}(\theta) = \hat{M}_{N_G}(\theta_{N_G}) \dots \hat{M}_1(\theta_1), \quad (\text{C1})$$

in which θ_i is the vector of parameters used by gate i , and then evaluates the cost function Eq. (23). This enables us to obtain analytic gradients of the cost function also as elements of products of matrices. We optimize the cost function using the BFGS method, and allow for multiple batches of input parameters with random variations added to avoid local minima. Additionally, as noted above, all of the isometries that result from the use of a real-valued quantum embedding map will be real, and so we can restrict our attention to real-valued gates. Hence, in our implementation, we parametrize single-qubit gates as y rotations

$$\hat{R}_y(\theta) \equiv \begin{pmatrix} \cos \frac{\theta}{2} & -\sin \frac{\theta}{2} \\ \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix}, \quad (\text{C2})$$

which relate to the gates in Eq. (22) as $\hat{R}_y(\theta) = \hat{U}_3(\theta, 0, 0)$, and CNOTs for the entangling gates. While we have made the above gate choices for use in this paper, we stress that

our methods apply to any other choice of single-qubit and entangling gates. While there is no guarantee that there are not operations with fewer entangling gates that could be found using complex-valued gates, we find that the reduction in the number of parameters when using real gates significantly improves the optimization time.

We utilize sparse representations of the individual parameterized gate elements $\mathbb{M}_a(\theta_a)$, but not for the full unitary $\hat{U}(\theta)$, as the latter is not guaranteed to be sparse. The matrix product in Eq. (C1) and the related expression for gradients dominates the computational scaling of the algorithm, which is $\mathcal{O}(\chi^3)$ in the bond dimension χ . This is the same asymptotic scaling with χ appearing in the DMRG-type optimization of the MPS network itself, described in Sec. IV A.

The final optimization we have included is to introduce longer gate sequence “motifs” into the optimization alongside the native entangling gates. In particular, the two motifs we have utilized in our work are a two-qubit rotation gate

$$\hat{S}(\theta, \theta') = \begin{pmatrix} \cos\left(\frac{\theta-\theta'}{2}\right) & 0 & 0 & \sin\left(\frac{\theta-\theta'}{2}\right) \\ 0 & \cos\left(\frac{\theta+\theta'}{2}\right) & \sin\left(\frac{\theta+\theta'}{2}\right) & 0 \\ 0 & -\sin\left(\frac{\theta+\theta'}{2}\right) & \cos\left(\frac{\theta+\theta'}{2}\right) & 0 \\ -\sin\left(\frac{\theta-\theta'}{2}\right) & 0 & 0 & \cos\left(\frac{\theta-\theta'}{2}\right) \end{pmatrix}, \quad (\text{C3})$$

which is allowed between any two qubits that have CNOT connectivity, and a version of the $\hat{S}$ gate we call $\hat{F}$ that is controlled on a third qubit. We find that the former gate can be compiled using two CNOTs using the ansatz sequence shown in Eq. (C4)

$$(\text{C4})$$

and the latter gate with control on qubit c and the operation $\hat{S}$ applied to qubits q_1 and q_2 can be constructed using

$$\begin{aligned} \hat{F}_{c;q_1q_2}(\theta, \theta') &= \text{CNOT}(c, q_2) \text{CNOT}(c, q_1) \hat{S}_{q_1q_2}\left(-\frac{\theta}{2}, -\frac{\theta'}{2}\right) \\ &\times \text{CNOT}(c, q_2) \text{CNOT}(c, q_1) \hat{S}_{q_1q_2}\left(\frac{\theta}{2}, \frac{\theta'}{2}\right). \end{aligned} \quad (\text{C5})$$

Hence, $\hat{S}$ gates require 2 CNOTs for compilation and $\hat{F}$ gates require 8 CNOTs for compilation. Both gates were identified from experiments with the greedy optimization procedure outlined above using only CNOTs, and their direct inclusion into the optimization enables more rapid convergence. As these gates require multiple entangling gates, it is useful to introduce a heuristic penalty function h into the cost function for ordering the next priority queue to ensure that they are not chosen over shorter gates with a similar cost function. Such a penalty function was advocated in Ref. [69], and could also be used to account for, e.g., hardware-dependent noise [79]. The choice of this penalty function will be problem-specific, and finding ways for optimizing it in a data-driven fashion for problems of interest is an intriguing area for further research.

FIG. 18. Example of greedy compilation procedure. Example gates represented as circuits and matrix plots resulting from applying the greedy compilation procedure to the isometry shown in the upper left (same isometry as in the right panel of Fig. 3). The starting ansatz is a single-qubit rotation on each qubit, given in the top center of the figure. The next row down shows the gates resulting from adding a single entangling gate to this ansatz, ordered left to right by their cost functions $\mathcal{C}$. A constant penalty 0.6 is added to the cost function for use of a $\hat{\mathcal{F}}$ gate in ordering the priority queue, resulting in the given ordering. The gates indicated by green lines denote those passed to the next level of optimization. This procedure terminates in the gate shown at the bottom of the figure with the given cost function tolerance of 5×10^{-4} .

We also note that the use of multiqubit controlled gates is penalized through the choice of the cost function Eq. (B7) for choosing the permutation to diagonal gauge; the choice of a cost function of 4 or 8 for a gate requiring two and three bit flips is in rough accordance with the number of CNOTs required for $\hat{\mathcal{S}}$ and $\hat{\mathcal{F}}$, respectively.

An example application of this procedure to the isometry shown in the right panel of Fig. 3 is given in Fig. 18. Here, we give cost function penalties of 0.6 and 0.2 for $\hat{\mathcal{F}}$ and $\hat{\mathcal{S}}$ gates, respectively, use a cost function tolerance of 5×10^{-4} , and keep the four lowest cost gates to generate the priority queue from the first optimization and the two lowest-cost gates on subsequent optimizations. The successive rows show the optimized gates resulting from adding a single entangling gate to the ansatz resulting from the last round of optimization, starting with a single-qubit rotation on each qubit (top center). The green lines show the gates which are kept to form the new priority queue. Here and throughout, the quantum circuits are ordered with the physical (i.e. readout) qubit on the top line and the ancilla qubits in increasing order on lower lines. Following an optimization in which $\hat{\mathcal{S}}$ or $\hat{\mathcal{F}}$ gates may be used, the “raw” circuit containing these parameterized gates is then compiled into CNOTs using Eqs. (C4) and (C5), products of single-qubit rotations are collected together, and then optimization passes are run to determine if single-qubit gates with rotations smaller than a certain threshold can be removed without affecting the cost function. We note that no cost function penalty is applied when an $\hat{\mathcal{S}}$ or $\hat{\mathcal{F}}$ gate brings the cost function below its desired tolerance, as in the last step of the optimization shown in Fig. 18, but is only used

for ordering the priority queue when no gates meet the cost function tolerance.

APPENDIX D: EXPLICIT CONSTRUCTION OF EXACTLY SOLVABLE MPS MODEL

In this section, we detail the transformation of the MPS representation of the exactly solvable model Eq. (24)–(26) to left-canonical form, defining a sequential preparation scheme. We then further detail a “by hand” compilation of a particular unitary completion of the resulting isometries as a benchmark for our automated procedure. For simplicity of exposition, we will take all phases $\phi_j = 0$, though we will relax this condition shortly. Performing the QR decomposition on the first tensor, we find

$$A^{[0]} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{p_0} \end{pmatrix} \quad (D1)$$

$$\rightarrow QR = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{\frac{p_0}{|p_0|}} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{p_0} \end{pmatrix}, \quad (D2)$$

$$\Rightarrow L^{[0]} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{\frac{p_0}{|p_0|}} \end{pmatrix}, \quad (D3)$$

$$A_{00}^{[1]0} = 1, \quad A_{01}^{[1]1} = \sqrt{p_1}, \quad A_{11}^{[2]0} = \sqrt{p_0}. \quad (D4)$$

Reshaping the second tensor and decomposing, we find

$$A_{(\alpha i)\beta}^{[1]} = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{p_1} \\ 0 & \sqrt{p_0} \\ 0 & 0 \end{pmatrix} \quad (\text{D5})$$

$$\rightarrow QR = \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{\frac{p_1}{p_0+p_1}} \\ 0 & \sqrt{\frac{p_0}{p_0+p_1}} \\ 0 & 0 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{p_0+p_1} \end{pmatrix}, \quad (\text{D6})$$

$$\Rightarrow L_{00}^{[1]0} = 1, \quad L_{01}^{[1]1} = \sqrt{\frac{p_1}{p_0+p_1}}, \quad L_{11}^{[1]0} = \sqrt{\frac{p_0}{p_0+p_1}}. \quad (\text{D7})$$

This generalizes to

$$L_{00}^{[j]0} = 1, \quad L_{01}^{[j]1} = \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}, \quad L_{11}^{[j]0} = \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}, \quad (\text{D8})$$

with the last tensor being

$$L_{10}^{[N-1]0} = \sqrt{\sum_{i < N-1} p_i}, \quad L_{00}^{[N-1]1} = \sqrt{p_{N-1}}. \quad (\text{D9})$$

We now take these left-canonical tensors and reshape them to correspond to isometries acting on a single physical qubit $|i_q\rangle$ and an ancilla qubit $|\alpha_a\rangle$. We start from the N^{th} tensor, where both the qubit and ancilla are in the state 0. The isometry is

$$\hat{L}^{[N-1]} = (\sqrt{1-p_{N-1}}|1_a 0_q\rangle + \sqrt{p_{N-1}}|0_a 1_q\rangle)(\langle 0_a 0_q|. \quad (\text{D10})$$

Following this, the physical qubit can be measured in the computational basis and its outcome (classically) stored, and then the physical qubit is returned to the state $|0_q\rangle$. We then repeat this procedure of acting with isometries, measuring the physical qubit, classically recording its output, and returning the physical qubit to 0, using the isometries

$$\hat{L}^{[j]} = |0_a 0_q\rangle\langle 0_a 0_q| + \left(\sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}|0_a 1_q\rangle + \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}|1_a 0_q\rangle \right) \langle 1_a 0_q|. \quad (\text{D11})$$

We note that Eq. (D11) also holds for the final site, $j = 0$, and produces an unentangled ancilla in the state $|0_a\rangle$. With these operators in hand, we can reinsert the arbitrary phases on the elements resulting in the state $|1_q\rangle$, yielding

$$\hat{L}^{[N-1]} = (\sqrt{1-p_{N-1}}|1_a 0_q\rangle + e^{i\phi_{N-1}}\sqrt{p_{N-1}}|0_a 1_q\rangle)(\langle 0_a 0_q|, \quad (\text{D12})$$

$$\hat{L}^{[j]} = |0_a 0_q\rangle\langle 0_a 0_q| + \left(e^{i\phi_j} \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}|0_a 1_q\rangle + \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}|1_a 0_q\rangle \right) \langle 1_a 0_q|. \quad (\text{D13})$$

FIG. 19. Circuit decompositions for $\hat{U}^{[j]}$ in Eq. (D14) (top) and $U^{[N-1]}$ in Eq. (D16) (bottom) based on a gateset of single qubit rotations and CNOTs. In both diagrams, the upper line is the physical (sampled) qubit and the lower line is the ancilla.

We note that there is a “natural” unitary completion of the operators in Eq. (D13) given by

$$\begin{aligned} \hat{U}^{[j]} &= |0_a 0_q\rangle\langle 0_a 0_q| + |1_a 1_q\rangle\langle 1_a 1_q| \\ &+ \left(e^{i\phi_j} \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}|0_a 1_q\rangle + \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}|1_a 0_q\rangle \right) \langle 1_a 0_q| \\ &+ \left(\sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}|0_a 1_q\rangle - e^{-i\phi_j} \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}|1_a 0_q\rangle \right) \langle 0_a 1_q|, \end{aligned} \quad (\text{D14})$$

in which the state $|1_a 1_q\rangle$ —that is never populated under ideal operation—is left unchanged and the action on the—also ideally unpopulated—state $|0_a 1_q\rangle$ is determined by orthogonality. Written in the basis representation $\{|0_a 0_q\rangle, |0_a 1_q\rangle, |1_a 0_q\rangle, |1_a 1_q\rangle\}$, we find

$$[\hat{U}^{[j]}] = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta_j & e^{i\phi_j} \sin \theta_j & 0 \\ 0 & -e^{-i\phi_j} \sin \theta_j & \cos \theta_j & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (\text{D15})$$

in which $\cos \theta_j = \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}}$ and so $\sin \theta_j = \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}}$. This gate has a natural interpretation as a rotation within the subspace of a single quantum of excitation shared between the qubit and ancilla, with the rotation angle set by the classical data vector probabilities [for $p_j \rightarrow 0$, $\theta_j \rightarrow 0$ and Eq. (D15) becomes the identity]. An analogous unitary completion for the isometry $\hat{L}^{[N-1]}$ is given by

$$\begin{aligned} \hat{U}^{[N-1]} &= (\cos \theta_{N-1}|1_a 0_q\rangle + e^{i\phi_{N-1}} \sin \theta_{N-1}|0_a 1_q\rangle)(\langle 0_a 0_q| \\ &+ (-e^{-i\phi_{N-1}} \sin \theta_{N-1}|1_a 0_q\rangle + \cos \theta_{N-1}|0_a 1_q\rangle)(\langle 0_a 1_q| \\ &+ (e^{i\phi_{N-1}} \sin \theta_{N-1}|1_a 1_q\rangle + \cos \theta_{N-1}|0_a 0_q\rangle)(\langle 1_a 0_q| \\ &+ (-e^{-i\phi_{N-1}} \sin \theta_{N-1}|0_a 0_q\rangle + \cos \theta_{N-1}|1_a 1_q\rangle)(\langle 1_a 0_q|. \end{aligned} \quad (\text{D16})$$

From a gate-based perspective, the operators in Eqs. (D15) with $\phi = -\pi/2$ are described by the Fermionic Simulation, or fSim(θ, φ) gate [80], with $\varphi = 0$ and $\theta = \theta_j$; this gate has been recently been demonstrated in gmon qubits [81]. Alternatively, they are a one-parameter generalization of the iSWAP gate [82], which can be compiled using the gate sequences shown in Fig. 19. We note that the unitary completion $\hat{U}^{[j]}$ at $\phi_j = 0$ is given by $\hat{S}(\theta_j, \theta_j)$ in the notation of Eq. (C3), and so for the gate set employed by the IBMQ processors, the

shortest decomposition for $\hat{U}^{[j]}$ is given by Eq. (C4). While in some alternative hardware platforms, such those employing tunable qubits [83], iSWAP gates can be implemented natively, partial iSWAPs still require decomposition. We also note that the operation Eq. (D15) is generated by the effective Hamiltonian

$$\hat{\mathbb{H}}^{[j]} = \theta_j (\hat{\sigma}_q^+ \hat{\sigma}_a^- + \hat{\sigma}_q^- \hat{\sigma}_a^+), \quad (\text{D17})$$

for “unit time” in the sense that

$$\hat{U}^{[j]} = \exp(-i\hat{\mathbb{H}}^{[j]}), \quad (\text{D18})$$

when $\phi_j = \pi/2$. This gate is readily achieved in trapped ion-based quantum computers using an equally weighted combination of XX and YY Mølmer-Sørensen gates [84], as well as a variety of other platforms implementing XY effective spin-spin interactions [85]. It is also interesting that the “data angle” θ_j has a natural interpretation as an ersatz “evolution time” in this perspective.

Before moving on from this exactly solvable example, we would like to point out how the freedom in representation of the bond basis manifests itself in this exactly solvable example. Namely, the predictions of the model Eq. (24)–(26) are unchanged if we reverse the roles of the $|0_a\rangle$ and $|1_a\rangle$ ancilla states in all but the first and last steps of preparation (using the unitary freedom exploited in the transformation to diagonal gauge discussed in Sec. III A). In this case, we have the isometries

$$\hat{L}^{[N-1]} = (\sqrt{1-p_{N-1}}|0_a 0_q\rangle + e^{i\phi_{N-1}}\sqrt{p_{N-1}}|1_a 1_q\rangle)\langle 0_a 0_q|, \quad (\text{D19})$$

$$\hat{L}^{[j]} = |1_a 0_q\rangle\langle 1_a 0_q| \quad (\text{D20})$$

$$+ \left(e^{i\phi_j} \sqrt{\frac{p_j}{\sum_{i \leq j} p_i}} |1_a 1_q\rangle + \sqrt{\frac{\sum_{i < j} p_i}{\sum_{i \leq j} p_i}} |0_a 0_q\rangle \right) \langle 0_a 0_q|. \quad (\text{D21})$$

The natural unitary completions of these isometries take the matrix representation

$$[\hat{U}^{[j]}] = \begin{pmatrix} \cos \theta_j & 0 & 0 & -e^{-i\phi_j} \sin \theta_j \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ e^{i\phi_j} \sin \theta_j & 0 & 0 & \cos \theta_j \end{pmatrix}, \quad (\text{D22})$$

and so are described by $\hat{S}(-\theta_j, \theta_j)$ at $\phi_j = 0$, and are generated by the effective Hamiltonian

$$\hat{\mathbb{H}} = \theta_j (\hat{\sigma}_q^+ \hat{\sigma}_a^+ + \hat{\sigma}_q^- \hat{\sigma}_a^-), \quad (\text{D23})$$

at $\phi = \pi/2$.

APPENDIX E: ISOMETRIES AND OPTIMIZED GATES FOR MNIST DATASET; $\chi = 8$

This Appendix contains plots of the isometries of a $\chi = 8$ MPS model for the MNIST example studied in Sec. V together with the compiled gate sequences obtained using the methods of Sec. III.

FIG. 20. Optimization for site 0. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 21. Optimization for site 1. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 22. Optimization for site 2. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 23. Optimization for site 3. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 24. Optimization for site 4. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 25. Optimization for site 5. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 26. Optimization for site 6. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 27. Optimization for site 7. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 28. Optimization for site 8. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 29. Optimization for site 9. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 30. Optimization for site 10. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 31. Optimization for site 11. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 32. Optimization for site 12. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 33. Optimization for site 13. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 34. Optimization for site 14. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 35. Optimization for site 15. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 36. Optimization for site 16. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 37. Optimization for site 17. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 38. Optimization for site 18. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 39. Optimization for site 19. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 40. Optimization for site 20. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 41. Optimization for site 21. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 42. Optimization for site 22. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 43. Optimization for site 23. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 44. Optimization for site 24. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 45. Optimization for site 25. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 46. Optimization for site 26. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 47. Optimization for site 27. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 48. Optimization for site 28. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 49. Optimization for site 29. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 50. Optimization for site 30. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 51. Optimization for site 31. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 52. Optimization for site 32. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 53. Optimization for site 33. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 54. Optimization for site 34. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 55. Optimization for site 35. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 56. Optimization for site 36. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 57. Optimization for site 37. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 58. Optimization for site 38. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 59. Optimization for site 39. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 60. Optimization for site 40. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 61. Optimization for site 41. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 62. Optimization for site 42. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 63. Optimization for site 43. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 64. Optimization for site 44. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 65. Optimization for site 45. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 66. Optimization for site 46. (a) Isometry. (b) Optimized gate. (c) Raw circuit from optimization. (d) Expanded and cleaned circuit from optimization.

FIG. 67. Optimization for site 47. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

FIG. 68. Optimization for site 48. (a) Isometry. (b) Optimized gate. (c) Circuit from optimization.

- [1] R. S. Smith, M. J. Curtis, and W. J. Zeng, A practical quantum instruction set architecture, [arXiv:1608.03355](https://arxiv.org/abs/1608.03355) [quant-ph].
- [2] D. S. Steiger, T. Häner, and M. Troyer, Projectq: An open source software framework for quantum computing, *Quantum* **2**, 49 (2018).
- [3] T. Häner, D. S. Steiger, K. Svore, and M. Troyer, A software methodology for compiling quantum programs, *Quantum Sci. Technol.* **3**, 020501 (2018).
- [4] G. Aleksandrowicz, T. Alexander, P. Barkoutsos, L. Bello, Y. Ben-Haim, D. Bucher, F. J. Cabrera-Hernández, J. Carballo-Franquis, A. Chen, C.-F. Chen, J. M. Chow, A. D. Córcoles-Gonzales, A. J. Cross, A. Cross, J. Cruz-Benito, C. Culver, S. D. L. P. González, E. D. L. Torre, D. Ding, E. Dumitrescu, I. Duran, P. Eendebak, M. Everitt, I. F. Sertage, A. Frisch, A. Fuhrer, J. Gambetta, B. G. Gago, J. Gomez-Mosquera, D. Greenberg, I. Hamamura, V. Havlicek, J. Hellmers, L. Herok, H. Horii, S. Hu, T. Imamichi, T. Itoko, A. Javadi-Abhari, N. Kanazawa, A. Karazeev, K. Krsulich, P. Liu, Y. Luh, Y. Maeng, M. Marques, F. J. Martín-Fernández, D. T. McClure, D. McKay, S. Meesala, A. Mezzacapo, N. Moll, D. M. Rodríguez, G. Nannicini, P. Nation, P. Ollitrault, L. J. O’Riordan, H. Paik, J. Pérez, A. Phan, M. Pistoia, V. Prutyantsev, M. Reuter, J. Rice, A. R. Davila, R. H. P. Rudy, M. Ryu, N. Sathaye, C. Schnabel, E. Schoute, K. Setia, Y. Shi, A. Silva, Y. Siraichi, S. Sivarajah, J. A. Smolin, M. Soeken, H. Takahashi, I. Tavernelli, C. Taylor, P. Taylor, K. Trabing, M. Treinish, W. Turner, D. Vogt-Lee, C. Vuillot, J. A. Wildstrom, J. Wilson, E. Winston, C. Wood, S. Wood, S. Wörner, I. Y. Akhalwaya, and C. Zoufal, Qiskit, <https://qiskit.org> (2020), [Online; accessed 17-February-2020].
- [5] R. LaRose, Overview and comparison of gate level quantum software platforms, *Quantum* **3**, 130 (2019).
- [6] Quantiki, List of QC Simulators, <https://www.quantiki.org/wiki/list-qc-simulators> (2020), [Online; accessed 17-February-2020].
- [7] E. T. Campbell, B. M. Terhal, and C. Vuillot, Roads towards fault-tolerant universal quantum computation, *Nature (London)* **549**, 172 (2017).
- [8] P. W. Shor, Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer, *SIAM Rev.* **41**, 303 (1999).
- [9] C. Gidney and M. Ekerå, How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits, [arXiv:1905.09749](https://arxiv.org/abs/1905.09749).
- [10] A. D. Córcoles, A. Kandala, A. Javadi-Abhari, D. T. McClure, A. W. Cross, K. Temme, P. D. Nation, M. Steffen, and J. M. Gambetta, Challenges and opportunities of near-term quantum computing systems, *Proc. IEEE* **108**, 1338 (2020).
- [11] J. Preskill, Quantum computing in the nisq era and beyond, *Quantum* **2**, 79 (2018).
- [12] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, Quantum machine learning, *Nature (London)* **549**, 195 (2017).
- [13] A. Perdomo-Ortiz, M. Benedetti, J. Realpe-Gómez, and R. Biswas, Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers, *Quantum Sci. Technol.* **3**, 030502 (2018).
- [14] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, Quantum machine learning: A classical perspective, *Proc. R. Soc. A* **474**, 20170551 (2018).
- [15] W. Huggins, P. Patil, B. Mitchell, K. B. Whaley, and E. M. Stoudenmire, Towards quantum machine learning with tensor networks, *Quantum Sci. Technol.* **4**, 024001 (2019).
- [16] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and J. I. Cirac, Expressive power of tensor-network factorizations for probabilistic modeling, with applications from hidden markov models to quantum machine learning, *Advances in Neural Information Processing Systems 32, Proceedings of the NeurIPS 2019 Conference* (2019).
- [17] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).
- [18] V. Havlíček, A. D. Córcoles, K. Temme, A. W. Harrow, A. Kandala, J. M. Chow, and J. M. Gambetta, Supervised learning with quantum-enhanced feature spaces, *Nature (London)* **567**, 209 (2019).
- [19] U. Schollwöck, The density-matrix renormalization group in the age of matrix product states, *Ann. Phys.* **326**, 96 (2011).
- [20] R. Orús, A practical introduction to tensor networks: Matrix product states and projected entangled pair states, *Ann. Phys.* **349**, 117 (2014).
- [21] R. Orús, Tensor networks for complex quantum systems, *Nat. Rev. Phys.* **1**, 538 (2019).
- [22] E. Grant, M. Benedetti, S. Cao, A. Hallam, J. Lockhart, V. Stojevic, A. G. Green, and S. Severini, Hierarchical quantum classifiers, *npj Quantum Inf.* **4**, 1 (2018).
- [23] Z.-L. Xiang, S. Ashhab, J. Q. You, and F. Nori, Hybrid quantum circuits: Superconducting circuits interacting with other quantum systems, *Rev. Mod. Phys.* **85**, 623 (2013).
- [24] M. Schuld, A. Bocharov, K. M. Svore, and N. Wiebe, Circuit-centric quantum classifiers, *Phys. Rev. A* **101**, 032308 (2020).
- [25] Y. LeCun, C. Cortes, and C. J. Burges, Mnist handwritten digit database, ATT Labs [Online]. Available: <http://yann.lecun.com/exdb/mnist> 2 (2010).
- [26] E. Stoudenmire and D. J. Schwab, Supervised learning with tensor networks, in *Advances in Neural Information Processing Systems* (Curran Associates, Inc., Red Hook, NY, 2016), pp. 4799–4807.
- [27] E. Farhi and H. Neven, Classification with quantum neural networks on near term processors, [arXiv:1802.06002](https://arxiv.org/abs/1802.06002).
- [28] M. Schuld and N. Killoran, Quantum Machine Learning in Feature Hilbert Spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [29] S. Lloyd, M. Schuld, A. Ijaz, J. Izaac, and N. Killoran, Quantum embeddings for machine learning, [arXiv:2001.03622](https://arxiv.org/abs/2001.03622).
- [30] M. Schuld, R. Sweke, and J. J. Meyer, Effect of data encoding on the expressive power of variational quantum-machine-learning models, *Phys. Rev. A* **103**, 032430 (2021).
- [31] R. LaRose and B. Coyle, Robust data encodings for quantum classifiers, *Phys. Rev. A* **102**, 032420 (2020).
- [32] S. R. White, Density Matrix Formulation for Quantum Renormalization Groups, *Phys. Rev. Lett.* **69**, 2863 (1992).
- [33] C. Schön, E. Solano, F. Verstraete, J. I. Cirac, and M. M. Wolf, Sequential Generation of Entangled Multiqubit States, *Phys. Rev. Lett.* **95**, 110503 (2005).
- [34] C. Schön, K. Hammerer, M. M. Wolf, J. I. Cirac, and E. Solano, Sequential generation of matrix-product states in cavity qed, *Phys. Rev. A* **75**, 032311 (2007).
- [35] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, Matrix product state representations, *Quantum Inf. Comput.* **7**, 401 (2007).

- [36] A. Cichocki, Tensor networks for big data analytics and large-scale optimization problems, [arXiv:1407.3124](#).
- [37] A. Cichocki, A.-H. Phan, Q. Zhao, N. Lee, I. V. Oseledets, M. Sugiyama, and D. Mandic, Tensor networks for dimensionality reduction and large-scale optimizations. part 2 applications and future perspectives, *Found. Trends Mach. Learn.* **9**, 431 (2017).
- [38] I. V. Oseledets, Tensor-train decomposition, *SIAM J. Sci. Comput.* **33**, 2295 (2011).
- [39] E. M. Stoudenmire, Learning relevant features of data with multi-scale tensor networks, *Quantum Sci. Technol.* **3**, 034003 (2018).
- [40] C. Guo, Z. Jie, W. Lu, and D. Poletti, Matrix product operators for sequence-to-sequence learning, *Phys. Rev. E* **98**, 042114 (2018).
- [41] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita, Reconstructing quantum states with generative models, *Nat. Mach. Intell.* **1**, 155 (2019).
- [42] G. Evenbly, Number-state preserving tensor networks as classifiers for supervised learning, [arXiv:1905.06352](#).
- [43] S. Klus and P. Gelß, Tensor-based algorithms for image classification, *Algorithms* **12**, 240 (2019).
- [44] S. Cheng, L. Wang, T. Xiang, and P. Zhang, Tree tensor networks for generative modeling, *Phys. Rev. B* **99**, 155131 (2019).
- [45] D. Liu, S.-J. Ran, P. Wittek, C. Peng, R. B. García, G. Su, and M. Lewenstein, Machine learning by unitary tensor network of hierarchical tree structure, *New J. Phys.* **21**, 073059 (2019).
- [46] I. Glasser, N. Pancotti, and J. I. Cirac, From probabilistic graphical models to generalized tensor networks for supervised learning, *IEEE Access* **8**, 68169 (2020).
- [47] M. Trenti, L. Sestini, A. Gianelle, D. Zuliani, T. Felser, D. Lucchesi, and S. Montangero, Quantum-inspired machine learning on high-energy physics data, [arXiv:2004.13747](#).
- [48] T.-D. Bradley, E. M. Stoudenmire, and J. Terilla, Modeling sequences with quantum states: A look under the hood, *Mach. Learn.: Sci. Technol.* **1**, 035008 (2020).
- [49] E. Gillman, D. C. Rose, and J. P. Garrahan, A tensor network approach to finite markov decision processes, [arXiv:2002.05185](#).
- [50] J. Miller, G. Rabusseau, and J. Terilla, Tensor networks for language modeling, [arXiv:2003.01039](#).
- [51] R. Selvan and E. B. Dam, Tensor networks for medical image classification, in *Proceedings of the Third Conference on Medical Imaging with Deep Learning* (2020).
- [52] J. Wang, C. Roberts, G. Vidal, and S. Leichenauer, Anomaly detection with tensor networks, [arXiv:2006.02516](#).
- [53] J. Reyes and M. Stoudenmire, A multi-scale tensor network architecture for classification and regression, [arXiv:2001.08286](#).
- [54] M. Lubasch, J. Joo, P. Moinier, M. Kiffner, and D. Jaksch, Variational quantum algorithms for nonlinear problems, *Phys. Rev. A* **101**, 010301(R) (2020).
- [55] A. S. Kardashin, A. V. Uvarov, and J. D. Biamonte, Quantum machine learning tensor network states, *Front. Phys.* **8**, 644 (2021).
- [56] A. V. Uvarov, A. S. Kardashin, and J. D. Biamonte, Machine learning phase transitions with a quantum processor, *Phys. Rev. A* **102**, 012415 (2020).
- [57] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, Unsupervised Generative Modeling Using Matrix Product States, *Phys. Rev. X* **8**, 031012 (2018).
- [58] S. Efthymiou, J. Hidary, and S. Leichenauer, Tensornetwork for machine learning, [arXiv:1906.06329](#).
- [59] Z.-Z. Sun, S.-J. Ran, and G. Su, Tangent-space gradient optimization of tensor network for machine learning, *Phys. Rev. E* **102**, 012152 (2020).
- [60] Y. Levine, O. Sharir, N. Cohen, and A. Shashua, Bridging Many-Body Quantum Physics and Deep Learning Via Tensor Networks, *Phys. Rev. Lett.* **122**, 065301 (2019).
- [61] J. Chen, S. Cheng, H. Xie, L. Wang, and T. Xiang, Equivalence of restricted boltzmann machines and tensor network states, *Phys. Rev. B* **97**, 085104 (2018).
- [62] I. Glasser, N. Pancotti, M. August, I. D. Rodriguez, and J. I. Cirac, Neural-Network Quantum States, String-Bond States, and Chiral Topological States, *Phys. Rev. X* **8**, 011006 (2018).
- [63] M. Foss-Feig, D. Hayes, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, J. M. Pino, and A. C. Potter, Holographic quantum algorithms for simulating correlated spin systems, [arXiv:2005.03023](#).
- [64] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson *et al.*, Demonstration of the qccd trapped-ion quantum computer architecture, [arXiv:2003.01293](#).
- [65] B. Coyle, D. Mills, V. Danos, and E. Kashefi, The born supremacy: Quantum advantage and training of an ising born machine, *npj Quantum Inf.* **6**, 1 (2020).
- [66] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi, Scaffcc: A framework for compilation and analysis of quantum computing programs, in *Proceedings of the 11th ACM Conference on Computing Frontiers* (Association for Computing Machinery, New York, NY, 2014), pp. 1–10.
- [67] M. Amy and V. Gheorghiu, staq-a full-stack quantum processing toolkit, *Quantum Sci. Technol.* **5**, 034016 (2020).
- [68] A. JavadiAbhari, S. Patil, D. Kudrow, J. Heckey, A. Lvov, F. T. Chong, and M. Martonosi, Scaffcc: Scalable compilation and analysis of quantum programs, *Parallel Comput.* **45**, 2 (2015).
- [69] M. G. Davis, E. Smith, A. Tudor, K. Sen, I. Siddiqi, and C. Iancu, Heuristics for quantum compiling with a continuous gate set, *Presented at the 3rd International Workshop on Quantum Compilation as part of the International Conference On Computer Aided Design 2019* (2019).
- [70] M. Soeken, F. Mozafari, B. Schmitt, and G. De Micheli, Compiling permutations for superconducting qpus, in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)* (IEEE, New York, NY, 2019) pp. 1349–1354.
- [71] R. Iten, R. Colbeck, I. Kukuljan, J. Home, and M. Christandl, Quantum circuits for isometries, *Phys. Rev. A* **93**, 032318 (2016).
- [72] F. Vatan and C. Williams, Optimal quantum circuits for general two-qubit gates, *Phys. Rev. A* **69**, 032315 (2004).
- [73] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. Jarrod Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Yu. Feng, E. W. Moore, J. Vand erPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python, *Nat. Methods* **17**, 261 (2020).
- [74] B. Efron, *The Jackknife, the Bootstrap and other Resampling Plans* (SIAM, Philadelphia, PA, 1982).

- [75] S. Bravyi, S. Sheldon, A. Kandala, D. C. McKay, and J. M. Gambetta, Mitigating measurement errors in multi-qubit experiments, [arXiv:2006.14044](#).
- [76] B. Nachman, M. Urbanek, W. A. de Jong, and C. W. Bauer, Unfolding quantum computer readout noise, [npj Quantum Inf.](#) **6**, 84 (2020).
- [77] J. Martyn, G. Vidal, C. Roberts, and S. Leichenauer, Entanglement and tensor networks for supervised image classification, [arXiv:2007.06082](#).
- [78] Y.-Y. Shi, L.-M. Duan, and G. Vidal, Classical simulation of quantum many-body systems with a tree tensor network, [Phys. Rev. A](#) **74**, 022320 (2006).
- [79] L. Cincio, K. Rudinger, M. Sarovar, and P. J. Coles, Machine learning of noise-resilient quantum circuits, [Phys. Rev. X Quantum](#) **2**, 010324 (2021).
- [80] I. D. Kivlichan, J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush, Quantum Simulation of Electronic Structure with Linear Depth and Connectivity, [Phys. Rev. Lett.](#) **120**, 110501 (2018).
- [81] B. Foxen, C. Neill, A. Dunsworth, P. Roushan, B. Chiaro, A. Megrant, J. Kelly, Z. Chen, K. Satzinger, R. Barends *et al.*, Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms, [Phys. Rev. Lett.](#) **125**, 120504 (2020).
- [82] N. Schuch and J. Siewert, Natural two-qubit gate for quantum computation using the xy interaction, [Phys. Rev. A](#) **67**, 032301 (2003).
- [83] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting qubits: Current state of play, [Annu. Rev. Condens. Matter Phys.](#) **11**, 369 (2020).
- [84] A. Sørensen and K. Mølmer, Entanglement and quantum computation with ions in thermal motion, [Phys. Rev. A](#) **62**, 022311 (2000).
- [85] T. Tanamoto, Y.-x. Liu, X. Hu, and F. Nori, Efficient Quantum Circuits for One-Way Quantum Computing, [Phys. Rev. Lett.](#) **102**, 100501 (2009).