

Enhancing Generative Models via Quantum Correlations

Xun Gao,¹ Eric R. Anschuetz,² Sheng-Tao Wang,^{3,1} J. Ignacio Cirac,^{4,5} and Mikhail D. Lukin¹

¹*Department of Physics, Harvard University, Cambridge, Massachusetts 02138, USA*

²*MIT Center for Theoretical Physics, 77 Massachusetts Avenue, Cambridge, Massachusetts 02139, USA*

³*QuEra Computing Inc., Boston, Massachusetts 02135, USA*

⁴*Max-Planck-Institut für Quantenoptik, Hans-Kopfermann-Str. 1, 85748 Garching, Germany*

⁵*Munich Center for Quantum Science and Technology (MCQST),*

Schellingstraße 4, D-80799 München, Germany



(Received 23 January 2021; revised 17 February 2022; accepted 17 March 2022; published 13 May 2022)

Generative modeling using samples drawn from the probability distribution constitutes a powerful approach for unsupervised machine learning. Quantum mechanical systems can produce probability distributions that exhibit quantum correlations which are difficult to capture using classical models. We show theoretically that such quantum-inspired correlations provide a powerful resource for generative modeling. In particular, we provide an unconditional proof of separation in expressive power between a class of widely used generative models, known as Bayesian networks, and its minimal quantum-inspired extension. We show that this expressivity enhancement is associated with quantum nonlocality and quantum contextuality. Furthermore, we numerically test this separation on standard machine-learning data sets and show that it holds for practical problems. The possibility of quantum-inspired enhancement demonstrated in this work not only sheds light on the design of useful quantum machine-learning protocols but also provides inspiration to draw on ideas from quantum foundations to improve purely classical algorithms.

DOI: 10.1103/PhysRevX.12.021037

Subject Areas: Interdisciplinary Physics

Quantum Physics, Quantum Information

I. INTRODUCTION

Over the past three decades, the field of machine learning has achieved remarkable success. A variety of powerful models and algorithms have been developed and deployed for broad applications ranging from computer vision and natural language processing to autonomous vehicles [1–3]. Unsupervised learning, involving the task of learning from unlabeled data sets, is among the frontier areas of machine-learning research. This task is typically much more challenging than supervised learning. The most common approach to tackle unsupervised learning problems is generative modeling, where one attempts to construct and train models with efficient representations for high-dimensional probability distributions. One of the most important aspects of any generative model is its expressive power, which, together with associated training algorithms, primarily determines the model performance. Models with high expressive power can capture complex correlations in the target probability distribution while upholding the

standard wisdom of Occam’s razor by keeping the structure simple (typically corresponding to a simple connectivity structure or limited number of parameters).

Quantum systems are known to produce complex probability distributions that are hard to capture with classical generative models [4–7]. For this reason, quantum models are believed to be more powerful in tackling unsupervised learning tasks. Consequently, over the past few years, quantum machine learning has emerged as a promising approach to enhance machine-learning performance [8–15]. However, apart from abstract computational complexity arguments [16–19], any potential quantum advantage in quantum machine-learning models and its physical origin is not well understood. Motivated by these considerations, in this work, we explore the role of quantum correlations associated with nonlocality and contextuality [20–23], both of which are known to be the key resource for quantum advantages in many quantum information processing tasks [24–28] and are expected to play a role in machine learning [11].

Specifically, we focus on a class of standard generative models, known as Bayesian networks, and show that quantum correlations can be used to achieve provable separation between such models and their minimal quantum-inspired extension described by a corresponding class of tensor networks. Focusing on sequential models,

Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International license](#). Further distribution of this work must maintain attribution to the author(s) and the published article’s title, journal citation, and DOI.

we compare subclasses of Bayesian networks with the corresponding 1D tensor networks described by matrix product states (MPSs) and show that MPSs feature more expressive power compared to traditional machine-learning models [29–34]. Since the 1D models can be efficiently evaluated on a classical computer, we also numerically test the models on real-world data sets and find an improvement in generative modeling using MPSs. These results provide new insights into the power of MPS-based machine-learning algorithms and open a fruitful direction to use ideas from quantum foundations to design new quantum and quantum-inspired machine-learning algorithms.

Our paper is organized as follows. In the next section, we provide an outline of the main results and discuss their implications. In Sec. III, we review Bayesian networks and their quantum circuit interpretation, and introduce our minimal quantum extension of Bayesian networks and its relation with tensor networks. In Sec. IV, we prove separations in expressivity between the two classes of models in learning sequential data sets. In Sec. V, we give numerical evidence that this separation often holds not only in theory but also in practice by showing separations on a variety of standard machine-learning data sets. Finally, in Sec. VI, we discuss the implications of our results and consider future lines of research.

II. SUMMARY OF RESULTS AND THEIR IMPLICATIONS

Bayesian networks, associated with a class of generative models based on directed graphs, have a wide range of applications [35]. Probability distributions described by Bayesian networks are known to have an equivalent formulation in the computational basis measurements of a class of quantum circuits known as Bayesian quantum circuits (see Fig. 1 and Ref. [36]). By extending this class to allow local measurement beyond the computational basis, we define a class of quantum-inspired models dubbed BBQCs, which are a special class of tensor networks that inherit the graph structure of their corresponding Bayesian networks.

In this work, we construct BBQCs that have unconditional expressivity separations compared to their classical counterparts, i.e., Bayesian networks on the same directed graphs. Instead of requiring an exact representation, we relax the comparison criterion to allow for any finite error in the forward and backward Kullback-Leibler (KL) divergence. This is equivalent to the condition

$$q(\mathbf{x}) = 0 \Leftrightarrow p(\mathbf{x}) = 0, \quad \forall \mathbf{x}. \quad (1)$$

where p and q are the two comparison distributions. However, this error model is still not practical enough; for instance, when $q(\mathbf{x}) = 0$ and $p(\mathbf{x})$ is very small, the exact KL divergence is infinite. In this paper, we adopt this error model to obtain rigorous proofs, but we show numerically that there exists a finite separation in KL

divergence even when, in practical training, $q(\mathbf{x})$ does not have exact zero probabilities for any \mathbf{x} . KL divergence is a widely used error model in unsupervised machine learning.

As a toy example, we first analyze the implications of quantum nonlocality for a so-called k -gram model, a very successful Bayesian network model used in natural language processing [see Fig. 1(a)] [37]. In practice, k is limited to be a constant since the number of model parameters, and hence the time and space resources, grow exponentially with k .

We introduce a basis-enhanced 2-gram model, shown in Fig. 2 (where the left is the BBQC and the right is the corresponding Bayesian network) and prove that any k -gram model for $k < \lfloor (n-1)/2 \rfloor$ cannot approximate its probability distribution under finite KL divergence. The proof is inspired by the mathematical structure of the nonlocal correlations present in measuring a Greenberger–Horne–Zeilinger state (GHZ state) that cannot be described by local hidden variable models [38]. We extend this argument to a cluster state where the qubits are measured either in the X or Y basis. This state can be represented by a basis-enhanced 2-gram model but not a local hidden variable model; thus, basis-enhanced 2-gram models exhibit correlations that share the mathematical structure of the quantum nonlocality present in the systems. For the corresponding k -gram model, however, the conditional probability distribution factorizes and can be described by a local hidden variable model. This result is summarized as the following theorem:

Theorem 1. (k -gram models and quantum nonlocality) There exists a family of basis-enhanced 2-gram models with generated probability distribution q such that any classical k -gram models with $k = o(n)$ (where n is the length of the 2-gram model) cannot approximate q to the error model in Eq. (1). This separation originates from quantum-inspired nonlocality present in the basis-enhanced model that is not present in the classical k -gram model.

We note that quantum-inspired correlations and quantum-inspired nonlocality (and later, contextuality) refer to the models that inherit the mathematical structure of the correlations present in quantum nonlocality and contextuality, but not the physical scenarios such as space-like separation; our proofs only rely on the mathematical properties of these quantum correlations. Since k -gram models can only capture local correlations, we also investigate a more expressive class of models, hidden Markov models [HMMs, shown in Fig. 1(b)], which are widely used in reinforcement learning and temporal pattern recognition. HMMs extend k -gram models by introducing hidden variables as memory to capture long-range correlations, and they are the most generic sequential generative models, including both feedforward and recurrent neural networks (given finite precision) as specific instances. We focus on the HMMs in the so-called translation form, with input and output regarded as original and target languages,

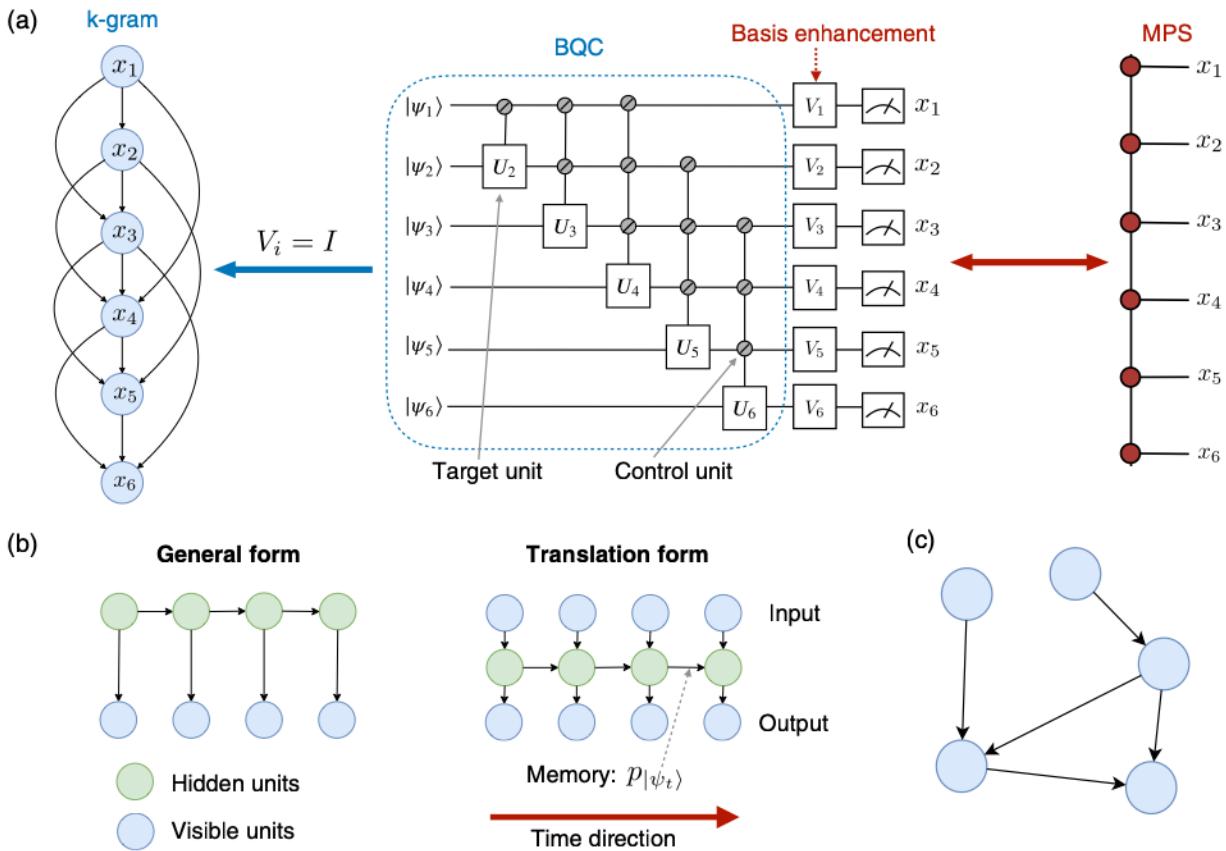


FIG. 1. Bayesian network and its quantum circuit. (a) k -gram model and its basis enhancement. The leftmost diagram is a 4-gram model as the transition probability takes the form $p(x_l|x_{l-1}, x_{l-2}, x_{l-3})$, involving four variables. All of the variables are visible. The middle diagram is a basis-enhanced Bayesian quantum circuit (BBQC). The blue dashed box contains a Bayesian quantum circuit (BQC) of the corresponding Bayesian network since measuring the output qubits in the computational basis results in the same probability distribution. The basic elements of BQCs are uniformly controlled gates for which control units and target units are labeled. The BBQC is a special case of tensor networks. In this case, the BBQC can be written as a MPS as shown in the rightmost diagram. (b) Hidden Markov model (left) and hidden Markov model for a translation problem (right). Visible variables and hidden variables are colored in blue and green, respectively. The top and bottom visible units store input and output, respectively. We adopt a dynamical point of view: The HMM can be understood as a measurement-driven evolution of probability distributions that encode the quantum states at the t th time step over the hidden variables in the t th virtual bond. (c) Example of a Bayesian network on a general directed graph. A more detailed discussion can be found in the Appendix B 3.

respectively, as shown in Fig. 1(b). Basis-enhanced versions of such HMMs correspond to a special instance of matrix product operators (MPOs). We show that MPOs exhibit correlations that share the mathematical structure of quantum contextuality present in the represented quantum systems. We directly show that the presence of this correlation gives rise to an expressivity separation between classical HMMs and their basis-enhanced counterparts. Specifically, we prove the following theorem:

Theorem 2. (Hidden Markov models and quantum contextuality) There exists a family of basis-enhanced 2-gram models, with a state space of dimensionality D , that cannot be approximated, in the sense of Eq. (1), by any classical hidden Markov models in the translation form [Fig. 1(b)], with a number of hidden units fewer than $D^{\Omega(\log D)}$. This separation originates from quantum-inspired

contextuality present in the basis-enhanced model that is not present in the classical HMM.

Here, the quantum-enhanced model is based on a 2-gram model [Fig. 5(c)], which is a special case of a HMM. The corresponding quantum circuit representation is shown in Figs. 5(a) and 5(b), with $D = 2^n$, where n is the number of qubits.

This result can be understood by considering the 1D structure of the models as a time dimension, as shown in Fig. 1(b). The state of the HMM (or the corresponding basis-enhanced 2-gram model) is encoded as a probability distribution $p_{|\psi_t\rangle}$ (quantum state $|\psi_t\rangle$) over the hidden states of the HMM (virtual bond of the MPO) at the t th time step. The number of hidden states (bond dimension) corresponds to the memory of the system, of which the logarithm is the number of bits (qubits) of memory required to store the

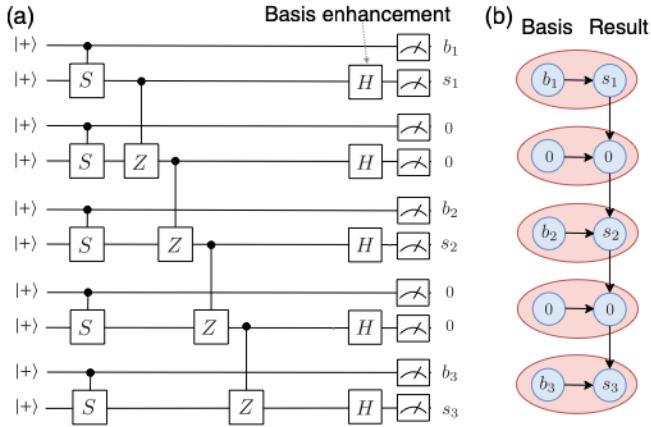


FIG. 2. (a) Example of a quantum 2-gram model. In this example, each pair of two qubits (counting from the top) corresponds to one variable. This BBQC can produce a $(2+3)$ -qubit cluster state and measure each qubit on the cluster state in either the X or Y basis. For each pair, the first qubit is used to choose the measurement basis; i.e., if it is measured to be 0, the second qubit in this pair will be measured in the X basis and otherwise in the $S^*XS = Y$ basis. The second qubit is used as the “physical” qubit in the cluster state. The cluster state can be extended to be arbitrarily long. Among all pairs, three pairs are measured respectively in the basis b_i with output results s_i . In between the three pairs, there is an odd number of ancillary pairs. These ancillary pairs are measured in the X basis with outputs postselected to be 0. There is a nonzero probability of getting a string that satisfies the constraint in Eq. (5). (b) Corresponding 2-gram model.

state of the system. The inputs and outputs are different measurement bases and measurement results, respectively. In order to simulate the quantum process, the HMM should have enough memory of the previous measurement basis and measurement results to predict future behavior. Within this picture, the translation form of HMMs is essentially equivalent to hidden variable models (also called noncontextual ontological models) [39–41]. Quantum contextuality formalizes the phenomenon that a measurement result of an observable should depend on which commuting observable set (known as a context) the observable belongs to in the given measurement scenario. However, since there are many different commuting sets that include this observable and there is no consistent preset assignments to all the observables (which is the key to contextuality), when it is measured, a hidden variable must memorize which context this observable belongs to in any given measurement scenario. A well-known example of contextuality is associated with the Mermin-Peres magic square [42,43]. Our proof strategy for Theorem 2 relies on showing that Mermin-Peres magic squares are very common in stabilizer states [44], and we use that feature to find a lower bound on the number of hidden states needed to accurately represent stabilizer measurements.

Finally, we evaluate the relative performance of BBQCs and Bayesian networks on standard machine-learning data sets. We focus on the relative performance of HMMs and their basis-enhanced counterparts, but here we use the general HMM graph structure in Fig. 1(b). As basis-enhanced HMMs are a special case of MPSs, we are able to evaluate the expressive performance of both HMMs and basis-enhanced HMMs efficiently on a classical computer. Specifically, we evaluate both models on the biofam data set [45,46], which is known to be well modeled by a simple 2-gram model. Additionally, we evaluate both models on the more difficult SPECT Heart and Promoter Gene Sequences data sets [47]. We find that the basis-enhanced HMM outperforms the HMM on both the training and testing data for both the SPECT Heart and Promoter Gene Sequences data sets and that it achieves comparable performance on the control biofam data set. These results are summarized in Fig. 7. In addition, we perform a likelihood-ratio test on the goodness of fits of the two models; this measures the statistical significance of the expressivity gap of the two models, accounting for the potential overfitting of the basis-enhanced model due to it having more parameters. We show that the improvement in KL divergence of the basis-enhanced HMM over the HMM is statistically significant on the SPECT Heart and Promoter Gene Sequences data sets to a confidence of greater than 5σ . These results are summarized in Fig. 8.

Our results have important implications for developing both classical and quantum machine-learning methods. Although the source of the advantage mechanisms described above is inspired by quantum correlations, for many classes of Bayesian networks—including k -gram and hidden Markov models—our extension still results in classical models, described by special cases of MPS or MPO that can be efficiently implemented on classical systems. In such cases, our results indicate that with a minimal computational overhead, one can obtain markedly improved modeling of data using novel quantum-inspired classical approaches. While a number of classical machine-learning techniques are already employing methods based on tensor networks [31,32,34,48], our results demonstrate that one could draw on ideas from quantum foundations to show unconditional efficiency separations for such novel classical models. Interestingly, our results also apply to models introduced by the classical machine-learning community compared with HMMs. One example is the observable-operator-model (OOM) [49] and its special case, the normed-observable-operator-model (NOOM) [50], which are generalization of HMMs to the complex number domain. Our methods also give a separation between these models and HMMs because of the normalization independence of the error model we considered.

Furthermore, for more complicated models, such as 2D Bayesian networks, where our extension cannot be efficiently implemented on a classical computer, our results

provide important insights into designing novel quantum machine learning algorithms. We emphasize that in contrast with the previously proposed quantum machine-learning models, which consider generic quantum circuits to provide quantum correlations, our approach makes use of the minimal extension of classical models. This is important since unstructured quantum circuits are not practical machine-learning models because of training difficulties associated with barren plateaus [51–55], poor local minima [56], and the no-free-lunch theorem [57]. Our results suggest that by restricting ourselves to minimal quantum extensions of classical machine-learning models, we may be able to sidestep these issues and maintain a quantum advantage over the corresponding classical model. In addition, this minimal approach allows us to understand the origin of the potential quantum advantage, which is essential for efficient design of new quantum machine-learning models. As an example, the technique of basis enhancement is also shown to be useful in other circuit-based quantum generative models in a subsequent work [58].

III. UNSUPERVISED GENERATIVE MODELING AND MINIMAL QUANTUM EXTENSIONS

Many unsupervised machine-learning tasks can be understood through a probabilistic lens. In this approach, the data \mathbf{x} [e.g., \mathbf{x} could be a vector (x_1, \dots, x_n) representing pixels of a handwritten digit] are regarded as being generated identically and independently from an unknown probability distribution p_D [1,3]. The task of unsupervised learning is to characterize some aspects of this distribution explicitly or implicitly. Generative models attempt to represent the entire probability distribution p_D approximately, thus providing an almost complete characterization of p_D .

Directly representing a probability distribution over (x_1, \dots, x_n) requires a number of parameters exponential in n . However, assuming some underlying structure on the distribution p_D , it is expected that only a polynomial number of parameters is sufficient to approximate p_D for most natural distributions. One can draw an analogy to quantum many-body physics: Physical states, which play the role of naturally occurring distribution p_D , typically require a polynomial number of parameters to represent, whereas generic states in the entire Hilbert space, like generic distributions with n variables, need an exponential number of parameters [59]. Graphical structures with a polynomial number of parameters often constitute efficient representations for generative models, similar to the representation of tensor networks in quantum many-body physics [60–64].

In what follows, we focus on a particular type of probabilistic graphical model, called Bayesian networks, and explore one minimal quantum—or, for simple enough networks, quantum-inspired—extension of this classical model. This allows us to understand the origin of the

underlying quantum enhancement or advantage, which sheds light on the design of quantum models. We emphasize that our approach is not limited to Bayesian networks and can be extended to other models.

A. Bayesian networks and language processing

Bayesian networks are a class of generative models that define a probability distribution through a directed acyclic graph in the following way [see example in Figs. 1(a) and 1(c)]: For each node x_i (associated with a random variable x_i), we assign a transition probability $p(x_i|\text{parents of } x_i)$, where the parents of x_i are nodes with edges directed towards x_i ; if there is no parent node for x_i , the transition probability reduces to the marginal probability $p(x_i)$; then, the product of these transition (marginal) probabilities,

$$p(x_1, \dots, x_n) = \prod_i p(x_i|\text{parents of } x_i), \quad (2)$$

is the final joint probability distribution.

Bayesian networks are useful in natural language processing as statistical language models. Roughly speaking, statistical language models are generative models for language, and they are used to generate a probability distribution of “meaningful” combinations of word sequences. A good design for the statistical language model is crucial to the performance of machine learning for natural language processing such as translation, speech recognition, and natural language generation [65].

Historically, prior to the rise of deep learning, one of the most commonly used statistical language models was the k -gram model [65], which is a Bayesian network on a 1D graphical structure with $k-1$ neighbors connected [see Fig. 1(a) for an example with $k=4$]. Despite their simplicity, certain types of generative neural networks can also be viewed as k -gram models, e.g., deep belief nets [66] even for $k=2$ (see Appendix A).

In order to capture more complex correlations, a more complex model, the HMM, with additional hidden nodes on a 1D graphical structure with nearest-neighbor connections, is introduced as shown in the translation form of Fig. 1(b). Note that k -gram models are a special case of HMMs (see Appendix A), as HMMs can store k -length correlations in the hidden variables. The graph structure of the HMM is easily generalized for translation problems. From a probabilistic point of view, a translation problem can be considered as a modeling problem for a conditional probability distribution $p(y|x)$ by a generative model, e.g., using the HMM shown in the translation form of Fig. 1(b), where \mathbf{x} (represented by the visible variables in the top row) is a “sentence” in the original language and \mathbf{y} (represented by the visible variables in the bottom row) corresponds to the translation in the target language with the conditional

probability $p(y|x)$. If the prior probability $p(x)$ can be captured by a HMM in the general form of Fig. 1(b), the joint probability distribution $p(x,y)$ could be viewed as a special case of the general form of Fig. 1(b).

B. Bayesian quantum circuits and minimal quantum extensions

The key to defining our quantum extension of Bayesian networks is the equivalence between Bayesian networks and a restricted class of quantum circuits, which we call Bayesian quantum circuits (see also Ref. [36]). BQCs are defined such that the probability distribution sampled from the quantum circuits, by measuring the visible qubits in the computational basis, is the same as the probability distribution defined by the corresponding Bayesian network. In addition, we define a minimal quantum extension of Bayesian networks, basis-enhanced Bayesian quantum circuits, by allowing the final measurements to be in an arbitrary local basis.

1. Bayesian quantum circuits

The building blocks of BQCs are uniformly controlled gates. A uniformly controlled gate is a generalization of a control- U gate, which consists of k control qubits and one target qubit [67]: If the control qubits are in the state $|x_1, \dots, x_k\rangle$, the target qubit will be applied by a unitary $U(x_1 \dots x_k)$ (i.e., a single qubit unitary determined by x_1, \dots, x_k). For convenience, we introduce the concepts of control units and target unit for a uniformly controlled gate as shown in Fig. 1(a).

Definition 1. (Bayesian quantum circuits) A Bayesian quantum circuit consists of a sequence of uniformly controlled gates followed by a measurement of a subset of the qubits in the computational basis, with the following restrictions:

- (i) Each uniformly controlled gate only targets a single qubit, reflecting the fact that there is only a single target variable in a transition probability in Bayesian networks.
- (ii) After being used as a control qubit, the qubit cannot be targeted by a uniformly controlled gate, reflecting the fact that Bayesian networks are defined on directed acyclic graphs.

The exact mapping between Bayesian networks and BQCs can be found in Appendix B 1. The implementation of an arbitrary uniformly controlled gate by elementary gates is, in general, not efficient since it typically consists of an exponential number of standard control gates [67]. However, for most relevant Bayesian networks, the transition probabilities only involve a few variables or are highly structured, which will make the corresponding uniformly controlled gates easy to implement. We further discuss the implementation of BQCs with multiqubit collective gates in Appendix B 2.

2. Basis-enhanced Bayesian quantum circuits

As defined above, BQCs can only produce distributions that correspond to Bayesian networks; thus, there is no quantum advantage in the expressivity of the model. In principle, there are many possible ways to generalize this model, such as violating the order requirement between target and control units or generalizing the uniformly controlled gates to more general gates. These generalizations will include universal quantum circuits and thus lose resemblance with classical Bayesian networks. To identify the differences between quantum models and their classical counterparts in terms of a quantum enhancement or advantage, we introduce a natural, minimal extension by allowing the measurements to be in other local bases beyond the computational basis. We call this basis-enhanced Bayesian quantum circuits. Note that the locality in the measurement basis is important; otherwise, the model will be as powerful as universal quantum circuits.

Definition 2. [Basis-enhanced Bayesian quantum circuits] A basis-enhanced Bayesian quantum circuit is a generalization of the Bayesian quantum circuit, where the measurements can be in any local basis beyond the computational basis.

This seemingly modest extension of classical Bayesian networks comes with considerable quantum advantages. For general underlying Bayesian networks, it can be shown that the quantum extension has an exponential improvement in expressive power compared to any “reasonable” classical generative models based on computational complexity assumptions (see Appendix B 3 for the rigorous proof). However, certain aspects of the complexity proof are unsatisfying. First, it relies on unproven computational complexity assumptions. Second, it does not provide physical insights and understanding on what gives rise to the purported quantum advantage. Therefore, in Sec. IV, we show explicitly that quantum correlations are the source of quantum enhancement and potential advantages for BBQCs. The unconditional separation between classical and quantum models based on quantum correlations is, however, more modest than the separation guaranteed by the complexity-theory-based arguments. The analysis can potentially be generalized to other models beyond Bayesian networks as discussed in Sec. VI.

3. BBQCs as tensor networks

Here, we remark that the quantum model BBQC is still a special case of tensor networks. We can use the k -gram model as an example, as shown in Fig. 1(a). Clearly, the BQC is a tensor network. Since the qubits are arranged on a line, we regard this tensor network as a MPS. The bond dimension is bounded by 2^{k-1} . The exponent is the maximum amount of information transmitted through qubits, and the unitary to change the measurement basis does not increase the bond dimension. Generally speaking, if the degree of the graph is bounded, the bond dimension around a

qubit in the tensor network is also bounded. Therefore, a BBQC can still be understood as a tensor network.

IV. PROVABLE EXPRESSIVITY SEPARATION THROUGH QUANTUM CORRELATIONS

To demonstrate how quantum correlations give rise to quantum enhancement, we compare the power of BQCs and BBQCs in generating sequential data. We show that, at least for some toy models, several fundamental nonclassical characteristics of quantum theory, i.e., nonlocality and contextuality, can be used as resources of quantum enhancement for generative models. At the same time, BBQCs are special classes of tensor networks. Our proof also demonstrates that ideas from quantum correlations can be used to show unconditional separation between purely classical models (k -gram versus MPS or HMM versus MPO).

A. Error models

Given a target probability distribution p_D and a distribution generated by a generative model p , one of the most commonly used cost functions to measure the effectiveness of p at modeling p_D is the forward KL divergence [3,68]:

$$D_{\text{KL}}(p_D \parallel p) = \sum_x p_D(x) \log \frac{p_D(x)}{p(x)}, \quad (3)$$

which is non-negative and lower bounded by 0 when $p = p_D$. Since KL divergence is asymmetric, one may also consider the reverse KL divergence, $D_{\text{KL}}(p \parallel p_D)$. The choice of forward versus reverse KL divergence when training unsupervised learning models reflects different priorities in the trained distribution [3].

To compare the expressive power of classical versus quantum models in generative modeling, we use the KL divergence to measure how effective classical models can generate a distribution originating from the corresponding minimally extended quantum model. In particular, we denote the probability distribution generated from BQCs as p and from BBQCs as q , and we investigate what the separation is between p and q in terms of expressive power.

The error model we use in the following theoretical analysis requires that both the forward KL divergence $D_{\text{KL}}(q \parallel p)$ and the reverse KL divergence $D_{\text{KL}}(p \parallel q)$ are finite. This is equivalent to Eq. (1). Thus, p approximating q under this error model is a weaker requirement than a small divergence of p from q . Nevertheless, we now show that quantum correlations, such as nonlocality and contextuality, give rise to quantum enhancement.

Before proceeding, we note that various other common error models can be considered. One widely used one, the multiplicative error, is a stronger requirement and a less realistic error model than the finiteness of KL divergence. This is used in our complexity-theory-based proof of the quantum advantage in Appendix B 3. In Sec. VI, we

discuss additional error models that are more realistic and more robust to small perturbations of model parameters. Relations among error models are explained in Appendix C.

B. Toy model: k -gram models and quantum nonlocality

In this section, we prove Theorem 1. The separation of 2 versus $O(n)$ between the basis-enhanced 2-gram model and classical k -gram model [Fig. 1(a)] is demonstrated through an example constructed from three-partite Bell tests of a GHZ state [38,69]. The GHZ state is embedded in an n -qubit 1D cluster state [70], such that measurement on $n - 3$ qubits in the X basis will produce a GHZ state (up to Pauli corrections according to the measurement results). A similar embedding was also used in Refs. [71,72]. The basis-enhanced 2-gram model is shown in Fig. 2(a), which can be verified directly to be a BBQC, where each variable corresponds to two qubits.

The measurement result b_i of the first qubit in the i th pair plays the role of choosing a measurement basis for the second qubit: $b_i = 0$ corresponds to measurement in the X basis, and $b_i = 1$ corresponds to measurement in the $Y = S^\dagger XS$ basis for the second qubit. All of the second qubits in each pair form a cluster state because they are connected through control-Z gates and the initial states are all $|+\rangle$. Suppose we choose three qubits to form the GHZ state and measure the remaining qubits according to Fig. 2(b). The resulting quantum state will be a GHZ state, where the probability to get b_i and s_i is

$$q(b_1 s_1 b_2 s_2 b_3 s_3 00\ldots 00) \propto 1 + i^{b_1+b_2+b_3} (-1)^{s_1+s_2+s_3}, \quad (4)$$

where b_i and s_i denote the measurement result from the first and second qubits of the i th pair. When $b_1 \oplus b_2 \oplus b_3 = 0$, the strings generated by this model with nonzero probability only contain b_i and s_i constrained by

$$i^{b_1+b_2+b_3} (-1)^{s_1+s_2+s_3} = 1. \quad (5)$$

It can be shown that any local hidden variable theory [i.e., $s_i = s_i(\lambda, b_i)$ does not depend on $b_{j \neq i}$, where λ is the hidden variable] cannot satisfy this equation [38,72].

We now prove that, for any $k < \lfloor (n - 1)/2 \rfloor$, a classical k -gram model cannot approximate the probability distribution q generated by the above BBQC up to the error defined in Eq. (1). We show this by reducing the classical model to a local hidden variable theory by fixing some variables, i.e., considering the conditional probability. When considering the full probability distribution, classical k -gram models are not strictly a “local” theory since there is information flow from the leftmost to the rightmost nodes even if k is a constant. There is causal influence between any pairs of nodes; i.e., a k -gram model can simulate the scenario that any node could communicate, though possibly only one way, to a node on the right. However, since

our error model in Eq. (1) only compares the support of the distributions, if one can show an expressivity separation in the conditional probability distributions, it directly implies a separation in the full probability distributions. For the conditional probability of the k -gram models, one can fix some variables, which can effectively cut the information flow. It turns out that the corresponding conditional probability of classical k -gram models is described by a local hidden variable theory:

$$p_C(b_1 s_1 b_2 s_2 b_3 s_3 | \text{other variables are } 0s) = \frac{f_1(s_1, b_1) f_2(s_2, b_2) f_3(s_3, b_3)}{\sum_{s_i, b_i} f_1(s_1, b_1) f_2(s_2, b_2) f_3(s_3, b_3)}, \quad (6)$$

where $f_i(s_i, b_i)$ is the product of the terms $p(b_{l+k-1} s_{l+k-1} | b_l s_l \dots b_{l+k-2} s_{l+k-2})$ involving s_i, b_i while setting other variables to 0. Because the three variables are chosen to be further than $2k$ apart, each product only involves one variable. We can thus normalize $f_i(s_i, b_i)$ to be $p(s_i | b_i, \lambda)$, where λ is determined by the measurement basis and results, as well as each term in the k -gram model, but λ does not depend on $b_{j \neq i}$. This shows that Eq. (6) can be described by a local hidden variable theory and thus completes the proof of Theorem 1.

In fact, our proof is very general. Using the chain rule of probability, one can show that if a conditional probability distribution factorizes (i.e., it is local), the model is a k -gram model with $k = O(1)$. This means any separation between the k -gram model and its basis-enhanced model must originate from the presence of correlations that share the mathematical structure of quantum nonlocality.

We note that the 2 vs $O(n)$ separation still holds under the error in Eq. (1), implying a separation also under the KL divergence. The circuit in Fig. 2 is essentially the same as the one used in Ref. [72] other than the boundary conditions. However, their result cannot be applied directly here since the k -gram model is not a constant-depth classical probabilistic circuit. Concretely, k does not correspond to the circuit depth, and the “light cone” scales with the system size even when k is small. On the other hand, hidden Markov models with bond dimension 6 could simulate this basis-enhanced 2-gram model. We give an explicit construction in Appendix F.

C. Hidden Markov models and quantum contextuality

We now study basis-enhanced HMMs [the translation form in Fig. 1(b)] in the context of translation problems. We find that any classical HMM requires $D^{\Omega(\log D)}$ hidden variables in order to approximate a basis-enhanced HMM with D hidden variables, under the error model of Eq. (1). The separation originates from quantum contextuality—in particular, our proof is constructed from the Mermin-Peres magic square [42,73].

Our approach is as follows. First, we discuss hidden variable theories—more precisely, noncontextual ontological theories [23,39,41]—and show that they are equivalent to classical hidden Markov models. Then, we give a lower bound on the number of ontological states needed to simulate Pauli measurements on stabilizer states using the Mermin-Peres magic square [73–75]. Finally, we discuss how basis-enhanced 2-gram models can efficiently simulate Pauli measurements on stabilizer states, proving our result.

1. Noncontextual ontological theories and hidden Markov models

First, we give a description of hidden variable theories (more precisely, noncontextual ontological theories) [23,39,41] in terms of hidden Markov models. At any moment, such an ontological theory is characterized by a state variable $\lambda_i \in \{\lambda_1, \lambda_2, \dots, \lambda_V\}$, which we assume completely determines the resulting distribution of the measurement outcomes of various observables. In particular, the model assumes that a quantum state $|\psi\rangle$ is encoded as a probability distribution over hidden variables as $p_{|\psi\rangle}(\lambda_i)$, where $\sum_i p_{|\psi\rangle}(\lambda_i) = 1$, and the measurement output from measuring an observable \hat{O} is described as

$$p_{|\psi\rangle}(y_i | \hat{O}) = \sum_i p(y_i | \lambda_i, \hat{O}) p_{|\psi\rangle}(\lambda_i), \quad (7)$$

where $p_{|\psi\rangle}(y_i | \hat{O})$ is the quantum mechanical measurement output probability for output y_i and $p(y_i | \lambda_i, \hat{O})$ is an indicator function independent of the quantum state $|\psi\rangle$. Below, we use the following notations as illustrated in Fig. 3(a). After a measurement of an observable \hat{O}_{x_i} from some restricted set of observables $\{\hat{O}_{x_i}\}$ labeled by x_i , and upon obtaining a measurement result y_i with probability $p(y_i | \lambda_i, x_i)$, there is also generally a transition probability to another state λ_j with probability $\Gamma_{(x_i, y_i)}(\lambda_j | \lambda_i)$, where $M = (x_i, y_i)$ characterizes the measurement and its result.

As discussed in Sec. III A, a HMM used for translation problems is a Bayesian network characterized at any moment by a hidden state λ'_i , with some input x_i , and it transitions to a new internal state λ'_j with probability $\alpha(\lambda'_j | \lambda'_i, x_j)$ and some probability $\beta(y_i | \lambda'_i)$ emitting a

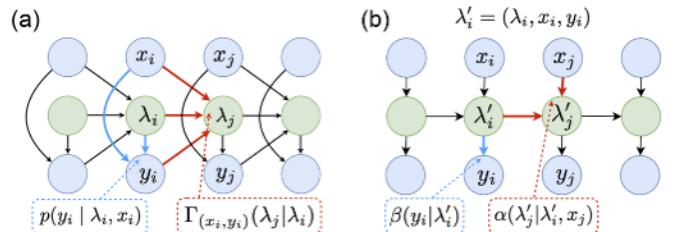


FIG. 3. (a) Ontological model in terms of Bayesian networks. (b) Standard hidden Markov model.

symbol y_i [see Fig. 3(b)]. We note that if we set $\lambda'_i = (\lambda_i, x_i, y_i)$, this is identical to a noncontextual ontological theory in the above paragraph.

2. Noncontextual ontological theories representing quantum states

Let us now consider how noncontextual ontological theories can simulate measurements on a quantum system. We follow Ref. [40] on the discussion of such ontological models, only keeping concepts that are relevant for our goal. We use Mermin-Peres magic squares explicitly to demonstrate the advantage of quantum contextuality.

A naive way to simulate a quantum system subject to sequential measurements is by recording each quantum state $|\psi_i\rangle$ that the system could generate as its own state variable λ_i . Though this encodes all information in a quantum theory, there is a large overhead in terms of the number of internal states λ_i needed, depending on which classes of circuits are modeled. We thus consider encodings that allow an internal state λ_i to be shared by different quantum states $|\psi\rangle$. In this case, each quantum state is encoded as a probability distribution over $\{\lambda_1, \lambda_2, \dots, \lambda_V\}$. Since we consider the error model in Eq. (1), we only need to consider if a measurement probability is zero or not, while the precise values of the probabilities are not important. Thus, a quantum state can be associated with a *support*

$$\text{supp}(|\psi\rangle) = \{\lambda_i | p_{|\psi\rangle}(\lambda_i) \neq 0\}, \quad (8)$$

which is the subset of internal states that the ontological theory could be in when representing the quantum state $|\psi\rangle$.

As illustrated in Fig. 1(b), we interpret the translation form of the HMM from a dynamical point of view. The state of the HMM is encoded as a probability distribution $p_{|\psi_t\rangle}$ over the hidden states of the HMM at the t th time step. The quantum state at time t , $|\psi_t\rangle$, depends on all the previous measurement outcomes, $M_t^{\text{past}} \equiv \{M_1, M_2, \dots, M_{t-1}\}$. Thus, in order to faithfully simulate the quantum process, the HMM should have enough memory about all previous measurement bases and outcomes to predict future behavior. The number of hidden states corresponds to the memory of the system. We could define the union of all the states at time t resulting from different measurement outcomes, $\{|\psi_t(M_t^{\text{past}})\rangle\}$, but there is ambiguity in setting the weights for different measurement histories. However, under the error model in Eq. (1), one only needs to be concerned with whether the probabilities are nonzero or not. Each measurement history can be associated with a support of hidden variables, i.e., $\text{supp}(|\psi_t(M_t^{\text{past}})\rangle)$, and the union over different histories can be defined as the union of the support spaces. We emphasize that this is well defined because there is no interference in the hidden Markov model; i.e., summation of different histories cannot be canceled.

Naively, one might believe that it is possible to encode $2^V - 1$ quantum states using only V ontological states since a set with V elements has $2^V - 1$ nontrivial subsets. However, in order for the ontological theory to make the same predictions as the quantum theory, there are restrictions on which subsets of hidden variables are used to label quantum states. For instance, if two states $|\psi_1\rangle$ and $|\psi_2\rangle$ are eigenstates with different eigenvalues of an allowed observable \hat{O} , then we must have $\text{supp}(|\psi_1\rangle) \cap \text{supp}(|\psi_2\rangle) = \emptyset$. As an example to illustrate this, suppose $\hat{O}|\psi_1\rangle = -|\psi_1\rangle$ and $\hat{O}|\psi_2\rangle = |\psi_2\rangle$, and there is at least one overlapping hidden variable in the support denoted as λ_i . Let us first assume the output is always $+1$ when measuring λ_i by \hat{O} , i.e., $p(+1|\lambda_i, \hat{O}) = 1$; then, there is a nonzero probability for both $|\psi_1\rangle$ and $|\psi_2\rangle$ to obtain the measurement result $+1$ according to Eq. (7), which contradicts the prediction from quantum measurements. Similarly, assuming the output of measuring λ_i by \hat{O} is always -1 or nonzero on both ± 1 also leads to the same contradiction.

A less trivial example is given by considering quantum contextuality, where the intersection among the supports of several quantum states should still be empty even if there is not a pair of states that are orthogonal. We now proceed to use the Mermin-Peres square to construct such triplets of states.

3. No common hidden variables in the Mermin-Peres magic square

In the following, we focus our attention on stabilizer states [44]. We first extend, in a more formal way, the discussion between memory and contextuality through the Mermin-Peres magic square example mentioned in Sec. II. Then, in Sec. IV C 4, we prove a lower bound on the number of hidden states required to simulate Pauli measurements on all stabilizer states.

Given three stabilizer states $|\psi_1\rangle$, $|\psi_2\rangle$, and $|\psi_3\rangle$, let $\{\hat{O}_{i1}\}$, $\{\hat{O}_{i2}\}$, and $\{\hat{O}_{i3}\}$ be their corresponding stabilizer groups. If

$$\{\hat{O}_{i1}\} \cup \{\hat{O}_{i2}\} \cup \{\hat{O}_{i3}\} \quad (9)$$

contains nine observables that form a Mermin-Peres square as shown in Table I, then we show, by contradiction, that the following must be true:

$$\text{supp}(|\psi_1\rangle) \cap \text{supp}(|\psi_2\rangle) \cap \text{supp}(|\psi_3\rangle) = \emptyset. \quad (10)$$

Concretely, we consider an example to illustrate this idea and leave the general proof to Appendix D. Let

$$\begin{aligned} |\psi_1\rangle &= |00\rangle, & |\psi_2\rangle &= |+\rangle, \\ |\psi_3\rangle &= \frac{|00\rangle + |01\rangle + |10\rangle - |11\rangle}{2}, \\ A &= Z_1, & a &= Z_2, & B &= X_2, & b &= X_1. \end{aligned} \quad (11)$$

TABLE I. Mermin-Peres magic square. The commutation relations among these operators are given in the following equations: $[A, a] = [B, b] = [A, B] = [a, b] = 0$ and $\{A, b\} = \{a, B\} = 0$.

$ \psi_1\rangle$	A	a	Aa	$+1$
$ \psi_2\rangle$	B	b	Bb	$+1$
$ \psi_3\rangle$	AB	ab	$ABab$	$+1$
	$+1$	$+1$	-1	

These stabilizer states and stabilizers form a Mermin-Peres magic square as in Table I. Here, we show that the intersection among the supports of the above three states must be empty. Suppose, instead, the intersection contains a hidden variable λ_i . We consider a measurement of the stabilizer $ABab = Y_1Y_2$. Since $|\psi_3\rangle$ is the eigenstate of $ABab$ with eigenvalue $+1$, we have that for any noncontextual ontological theory in state λ_i belonging to the support of $|\psi_3\rangle$,

$$p(y = +1|\hat{O}_y = Y_1Y_2, \lambda_i) = 1, \quad (12)$$

in order to make the same prediction with quantum mechanics and since $\sum_i p_{|\psi_3\rangle}(\lambda_i) = 1$. In addition, there exists a λ_j such that

$$\Gamma_{(Y_1Y_2,+1)}(\lambda_j|\lambda_i) > 0. \quad (13)$$

Since λ_i also belongs to the supports of $|\psi_1\rangle$ and $|\psi_2\rangle$, λ_j must also belong to the supports of

$$\begin{aligned} \frac{(1 + Y_1Y_2)}{2}|\psi_1\rangle &\propto \frac{1}{\sqrt{2}}(|00\rangle - |11\rangle), \\ \frac{(1 + Y_1Y_2)}{2}|\psi_2\rangle &\propto \frac{1}{\sqrt{2}}(|01\rangle + |10\rangle), \end{aligned} \quad (14)$$

which are the resulting states after measuring Y_1Y_2 on states $|\psi_1\rangle$ and $|\psi_2\rangle$ and getting the measurement result $+1$. However, these two states are orthogonal and thus cannot share a common λ_j as explained in Sec. IV C 2 (e.g., consider measuring Z_1Z_2 or X_1X_2). We thus arrive at a contradiction, and the three states cannot share a common hidden variable λ_i . It is straightforward to extend this example to more general stabilizers with the same commutation relations as those in Table I, which is detailed in Appendix D.

4. Bounding the efficiency of hidden Markov models

We now prove a lower bound, which has been used in Ref. [40], for the number of hidden states needed in the HMM. Denote S as the set of all the possible quantum states appearing in the quantum system that we want to simulate using a noncontextual ontological theory; in our case here, it is the set of all stabilizer states. Denote s as a subset of S such that

$$\bigcap_{|\psi\rangle \in s} \text{supp}(|\psi\rangle) \neq \emptyset. \quad (15)$$

Let

$$m = \max_s |s|. \quad (16)$$

Then, the number of state variables V needed in such an ontological theory in order to simulate the quantum system is not smaller than $|S|/m$. The reason for this is illustrated in Fig. 4.

The following two lemmas show that by allowing for all stabilizer states and all Pauli measurements, we have large $|S|$ and relatively small m . We upper bound m by using contextuality and proving the existence of the Mermin-Peres square.

Lemma 1. (Proposition 1 in Ref. [76]) The total number of stabilizer states is

$$|S| = 2^{n^2/2+o(n^2)}, \quad (17)$$

where n is the number of qubits.

Lemma 2. For any subset of stabilizer states s such that $|s| > 2^{n^2/4+7n/2}$, there exist three states such that some of their stabilizers form a Mermin-Peres magic square as shown in Table I. Therefore, $m \leq 2^{n^2/4+7n/2}$.

The idea of the proof is as follows: (1) Up to classical Clifford circuits (those only composed of CNOT and x), any stabilizer state can be written as a tensor product of a uniform superposition state of k qubits (with possibly nontrivial phases) and a product state of $n - k$ over the computational basis; (2) for any $2^{n^2/4+3n/2} \cdot 4^n \cdot 2$ states, we can always find one computational basis product state and $4^n \cdot 2$ other states that are superpositions over the same k qubits up to a Clifford circuit; (3) we can always find

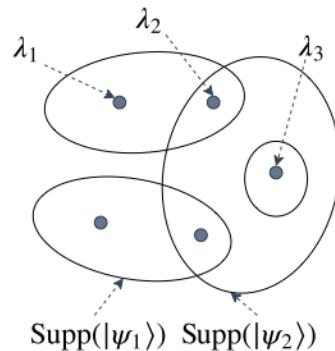


FIG. 4. Each circle corresponds to a support of a quantum state; there are thus $|S|$ circles. Each dot corresponds to a hidden variable. We denote the number of dots as V , and the number of circles that a dot belongs to is, at most, m . Imagine that we eliminate a dot and its associated circles each time. After V steps, all of the circles are eliminated. Since each time we can eliminate at most m circles, we have $Vm \geq |S|$.

2 standard graph states from these $4^n \cdot 2$ states up to single-qubit phase gates (which do not change the computational basis product state); (4) finally, the stabilizers of the computational basis product state and the 2 graph states can always form a Mermin-Peres square. Details of the proof are in Appendix D. From both lemmas, we have a lower bound for V :

$$V \geq \frac{|S|}{m} \geq 2^{(n^2/4)-O(n)}. \quad (18)$$

Following the discussion on the equivalence of noncontextual ontological theories and HMMs in Sec. IV C 1, the lower bound on V shows that any “translation-form” HMM simulating Pauli measurements on stabilizer states on n qubits require at least $2^{\Omega(n^2)}$ hidden states. In the following section, we show that a basis-enhanced 2-gram model, which is a special case of the translation-form HMM as in Fig. 1(b) (see Appendix B 1), only needs $2^{O(n)}$ internal states.

We have three remarks on the above proof. First, the idea does not need to be restricted to stabilizer states. Here, we use stabilizer states for the simplicity of illustrating the key idea, i.e., making use of the simplest example of contextuality, the Mermin-Peres square. Second, the above argument also works for HMMs without translational invariance. One can remove the requirement of translational invariance by just relabeling $\Gamma^{(t)}$ and $\text{supp}^{(t)}$ for the t th time step. Finally, as we show that HMMs are equivalent to noncontextual ontological models, our proof also implies that any such model that can express this language with $2^{o(n^2)}$ hidden states cannot be a noncontextual ontological model. This shows that this separation stems from the correlations of quantum contextuality present in the basis-enhanced hidden Markov model.

5. Basis-enhanced 2-gram model from stabilizer states

Here, we construct a basis-enhanced 2-gram model, as shown in Fig. 5, that simulates Pauli measurements on stabilizer states using $O(n)$ qubits—the underlying 2-gram model, therefore, has $D = 2^{O(n)}$ internal states. The input and output of the HMM correspond to the choices of different Pauli measurements and the measurement results, respectively, and the teleportation gadgets connect successive nodes in the 2-gram model. Together with the result in Sec. IV C 4 on the lower bound for the number of hidden variables required in a classical HMM to simulate the quantum process, $V \geq 2^{\Omega(n^2)} = D^{\Omega(\log D)}$, we have thus completed the proof of Theorem 2. Furthermore, instead of regarding it as a 1D model with a large bond dimension, we can also view it as a 2D model: In Fig. 5(a), if each qubit is regarded as a node in the Bayesian network, the corresponding directed graph is shown in Fig. 5(d).

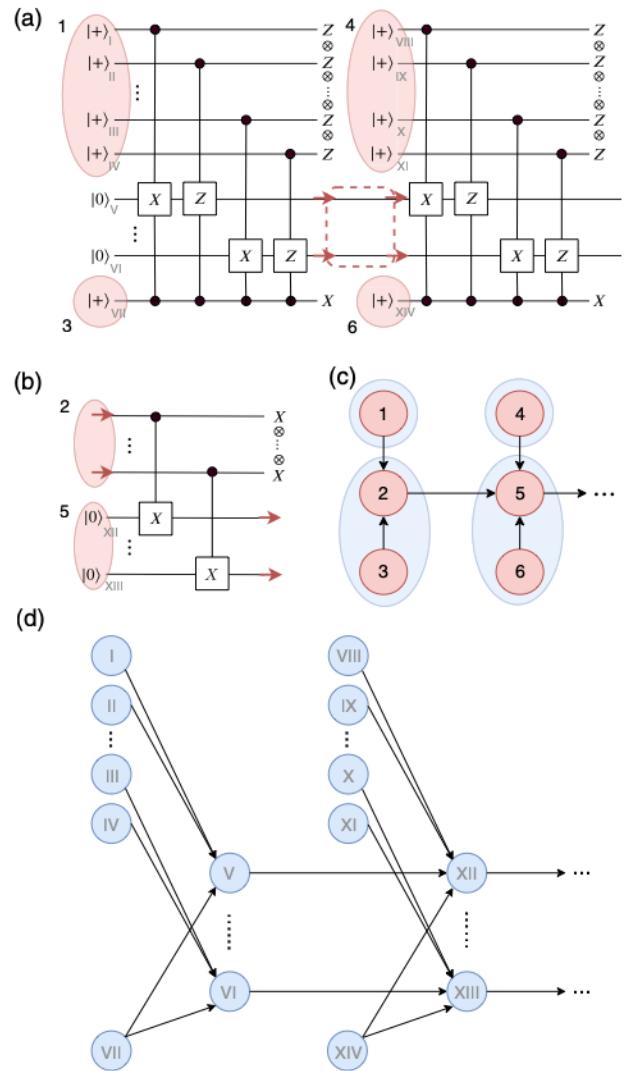


FIG. 5. (a) Quantum circuits for generating all of the possible stabilizer states by measurements. Circles 1 and 4 are superpositions of $2n$ -bit strings with uniform weights, of which measurement results in the Z basis determine which stabilizers are to be measured for the qubits initialized as $|0\rangle^{\otimes n}$ (those qubits in the middle). The measurement results in the X basis of circles 3 and 6 give the outcome of measuring the corresponding Pauli determined by the measurement result of circles 1 and 4, respectively. With all of the outcomes of circles 1,3,4,6, ..., the qubits in the middle can be any possible stabilizer state. (b) Teleportation gadget for the part in the red dashed box of diagram (a). The stabilizer to be measured in the next time step should be determined by measurement results of both 2 and 5 because the measurement result of 2 will sometimes produce some Pauli corrections. (c) Corresponding Bayesian network circuit of (a) after using (b). All of the red circles correspond to circles in (a) and (b) with the same number labels. We group $\{1\}$, $\{2, 3\}$ and $\{4\}$, $\{5, 6\}$ as blue circles to form a 2-gram model with input, which is a special case of the translation form of the HMM given in Fig. 1(b). (d) 2D arrangement of the Bayesian network for the circuit shown in (a). The correspondence between quantum circuits and circles is labeled by Roman numerals.

Before concluding this comparison between HMMs and basis-enhanced HMMs, we note that although this basis-enhanced circuit is constructed from stabilizer states, which can be efficiently simulated classically, this does not imply that a quantum computer is not useful in this case. In particular, one can consider continuous Pauli rotations instead of Clifford gates, which will be required in practice in order to train the model using a method such as gradient descent. Such quantum training algorithms for these models are analyzed in Appendix B 4. In general, such algorithms cannot be efficiently simulated classically. Furthermore, more general models than HMMs that are beyond one dimensional would need to sample from a quantum computer.

V. NUMERICAL TESTS ON REAL-WORLD DATA

In the previous section, we have proven theoretically that the quantum-inspired models we consider have more expressive power than the corresponding classical models. The sources of the quantum enhancement are quantum nonlocality and contextuality. In this section, we numerically test that the quantum models do indeed have better performance in practice. These numerical results primarily serve two purposes. First, they demonstrate that the quantum models actually have a potential advantage on real-world data. Second, they show that the separation is robust to more practical error models beyond the one used theoretically as in Eq. (1).

Concretely, we focus on classical HMMs and the quantum extension of 2-gram models introduced in Sec. IV C. As in most generative modeling tasks, the quantity of interest to evaluate the performance of the parametrized model p_{model} given a data set p_{data} is the forward KL divergence

$$D_{\text{KL}}(p_{\text{data}} || p_{\text{model}}) = \sum_{\mathbf{m} \in [1, \dots, M]^n} p_{\text{data}}(\mathbf{m}) \log \frac{p_{\text{data}}(\mathbf{m})}{p_{\text{model}}(\mathbf{m})}; \quad (19)$$

consistent with our convention used in previous sections, we let M denote the dimensionality of a given visible node in our model, and we let n be the number of visible nodes in the model. Since summing over the exponential number of terms in Eq. (19) is intractable in practice, we use the stochastic estimate of the KL divergence given in Eq. (B7).

A. Simulation of basis-enhanced 2-gram models

We now focus on (translationally invariant) classical hidden Markov models and basis-enhanced 2-gram models, both of which were introduced in Sec. IV C. Though in Sec. IV C we considered a specific translation task for the sake of our analysis, here we consider general basis-enhanced 2-gram models, with the parameters trained to represent some given data set. The general structure of the model we consider is given in Fig. 6.

Though basis-enhanced 2-gram models cannot directly be interpreted as classical Bayesian networks, they are still classically simulable using tensor networks when they have low bond dimension [32,48], making them a natural choice for numerical tests of our analysis [see Fig. 6(b)]. In particular, the direction of steepest descent of Eq. (B7) when varying a particular tensor U is given by its negative gradient with respect to the conjugate of the parameters [77]; that is, the direction of steepest ascent with respect to U takes the form

$$\frac{\partial \tilde{D}_{\text{KL}}}{\partial U} = \frac{2\partial_{\bar{U}} Z}{Z} - \frac{2}{K} \sum_{\mathbf{m} \in \text{data}} \frac{\partial_{\bar{U}} S_U(\mathbf{m})}{S_U(\mathbf{m})}, \quad (20)$$

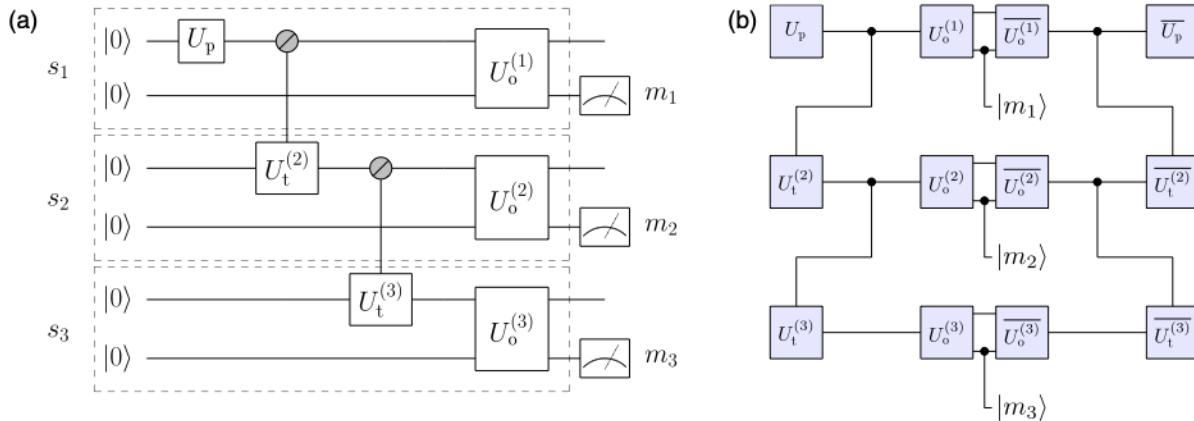


FIG. 6. (a) BBQC interpretation of a three-node 2-gram model, with states s_1, s_2, s_3 . An initial unitary U_p constructs the prior, and uniformly controlled unitaries U_t encode transitions between 2-gram states. Each measurement is in some standard basis encoding of the M -dimensional output space. To directly compare with classical HMMs, we only consider transfer unitaries on a k -dimensional subspace of each state (top wire), and we consider the distribution on an M subspace of each state (bottom wire). (b) Tensor network representation of the same quantum circuit.

where S is the unnormalized probability distribution given in Fig. 6(b) and Z is its normalization. When we perform the Riemannian descent algorithm described in Sec. V B, we optimize on the manifold of unitary matrices, and thus $Z = 1$. As we maintain translational invariance in our model, the total derivative with respect to some parameter \bar{U} is given by the sum of the variation over all equivalent tensors:

$$\frac{\partial \tilde{D}_{KL}}{\partial \bar{U}} = \sum_j \frac{\partial \tilde{D}_{KL}}{\partial \bar{U}^{(j)}}. \quad (21)$$

For completeness, we give examples of the tensor network representations of $\partial_{\bar{U}} S_U(\mathbf{m})$ and $\partial_{\bar{U}} Z$ in Appendix G. Since within one training minibatch many of the same tensors are contracted, in practice, we precompute intermediate tensor contraction results for each minibatch. For a basis-enhanced 2-gram model with bond dimension h , the classical runtime is $O(nh^3M)$ for computing the gradient with respect to the unitaries in the model. For comparison, a classical HMM trained using the Baum-Welch algorithm [78] takes time $O(nh(h + M))$ per training iteration.

B. Model training

In general, training Bayesian networks beyond tree graphs is hard [79,80]. We note that there exist many heuristic and approximate algorithms that work well in practice for training classical Bayesian networks [81], and we consider similar heuristics for BBQCs here, as described in more detail in Appendix B 4.

Since we focus on translationally invariant HMMs here, we can use the Baum-Welch algorithm to efficiently train the classical model [78]. Furthermore, as discussed in Sec. V A, computing the gradient of the loss function with respect to the parameters in the basis-enhanced 2-gram model is classically efficient using tensor networks for small bond dimension. However, naively performing gradient descent on the parameters of the model would generally violate unitarity constraints in the underlying quantum circuit model. Therefore, to optimize the unitaries used in the construction of the quantum model, we perform a variant of the Riemannian gradient descent algorithm introduced in Ref. [82].

Normally, in the gradient descent of some loss function $\mathcal{L}(\{A_i\})$ for complex-valued matrices A_i , one iteratively estimates the optimal A_i through the update rule [77]

$$A_i \rightarrow A_i - \alpha \frac{\partial \mathcal{L}}{\partial A_i}, \quad (22)$$

where α is the learning rate. In practice, keeping a moving average of previous gradient estimates smooths out stochastic fluctuations in estimates of $\partial \mathcal{L} / \partial A_i$; thus, we consider the momentum-based update rule [83]:

$$v_{A_i} \rightarrow \beta v_{A_i} + \alpha \frac{\partial \mathcal{L}}{\partial A_i}, \quad (23)$$

$$A_i \rightarrow A_i - v_{A_i}. \quad (24)$$

For unitary A_i , however—as in the case of quantum circuits—this procedure will generally yield nonunitary A_i . Therefore, we analytically calculate the direction of steepest descent in unitary space in terms of $\partial \mathcal{L} / \partial \bar{A}_i$ and perform parallel transport in that direction [82]. This leads to the update rule for a unitary matrix U :

$$U \rightarrow \exp \left(-\alpha \left(U \left(\frac{\partial \mathcal{L}}{\partial \bar{U}} \right)^\dagger - \frac{\partial \mathcal{L}}{\partial \bar{U}} U^\dagger \right) \right) U. \quad (25)$$

We modify the method in Ref. [82] slightly to allow for the momentum update rule of Eq. (24); namely, we use the update rule

$$U \rightarrow \exp \left(- (U v_U^\dagger - v_U U^\dagger) \right) U. \quad (26)$$

C. Model comparison on data sets

We test the performance of our implemented quantum extension of a 2-gram model on three data sets: the biofam (sequence length $n = 16$, output dimensionality $M = 8$) [45,46], Promoter Gene Sequences ($n = 57$, $M = 4$) [47], and SPECT Heart ($n = 23$, $M = 2$) [47] data sets. For all of our simulations, we use $\beta = 0.5$. For the biofam data set [45,46], we use $\alpha = 10^{-3}$, and for the Promoter Gene Sequences and SPECT Heart data sets [47], we use $\alpha = 10^{-2}$. For the biofam data set, we train for 75 epochs, and for the Promoter Gene Sequences and SPECT Heart data set, we train for 150 epochs. For all data sets, we estimate the gradient over a minibatch size of 8 training samples. The biofam data set tracks the family life of individuals from year to year (e.g., married, divorced, married with children) and is correlated from year to year. We expect it to be efficiently captured by a classical HMM due to the local nature of the data, and we use it as a control. The Promoter Gene Sequences data set consists of DNA sequences that encode promoters and nonpromoters; therefore, it has a less-obvious local structure. Finally, the SPECT Heart data set encodes binary feature vectors of heart images, with little to no local correlations.

To estimate the generalization performance of the models, we withhold a quarter of the data for testing the biofam and Promoter Gene Sequences data set, and use the standard SPECT Heart testing data set. Our results are summarized in Fig. 7, where we plot the stochastic estimate of the KL divergence \tilde{D}_{KL} (normalized by the sequence length) as a function of the local hidden dimension k . As we are interested in the optimal performance over all parameters to compare the expressive power of quantum-inspired

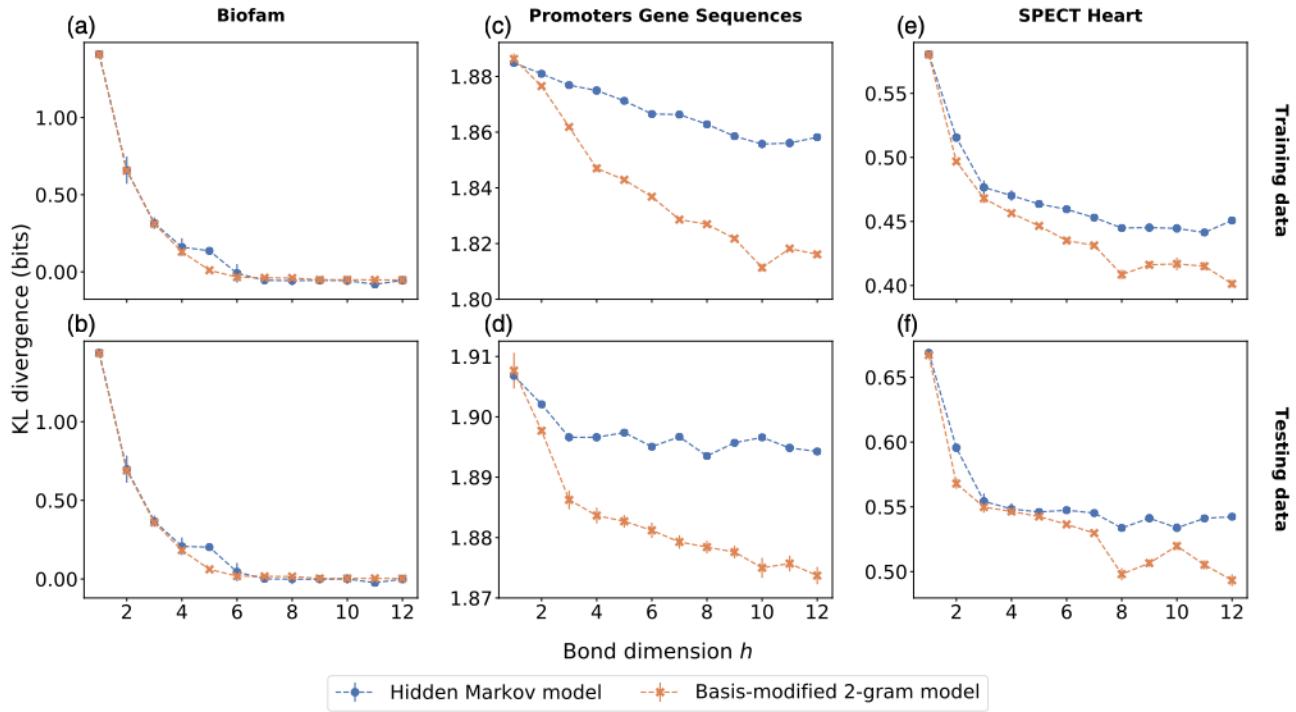


FIG. 7. Best performances over ten trials of the classical HMM (blue circles) and the basis-enhanced 2-gram model (orange crosses) on the (a,b) biofam, (c,d) Promoter Gene Sequences, and (e,f) SPECT Heart data sets. The first row plots the performance on the training data, and the second the performance on withheld testing data. The basis-enhanced 2-gram models have better performance than classical HMMs on the SPECT Heart and Promoter Gene Sequence data sets. Error bars denote 1 standard error of the mean over ten trials. Dashed lines are to aid the eye.

models versus classical models, we plot the minimum achieved loss over ten trials. In particular, for the Promoter Gene Sequences and SPECT Heart data sets, the basis-enhanced 2-gram model learns the distribution of samples more effectively and also generalizes more effectively than the classical HMM. As expected, both models perform equally well on the biofam data set since it has very local correlations. These results also demonstrate that for data sets that have no obvious local structure, quantum models tend to perform better, which is consistent with our theoretical analysis in Sec. IV C. Furthermore, we perform a likelihood-ratio test between the two models to measure the statistical significance of the improvement in performance, accounting for any potential overfitting due to the quantum model having more parameters than the classical model. Taking the null hypothesis that the optimal parameters of the basis-enhanced 2-gram model reduce the model to a classical hidden Markov model, using the observed difference in achieved KL divergence, we find that the null hypothesis can be rejected with 5σ confidence on the Promoter Gene Sequences and SPECT Heart data sets (see Fig. 8).

Interestingly, the performance separation between quantum and classical models persists even when considering the average performance over many runs. These results are summarized in Appendix G. With these specific examples, the sets of numerical results suggest that quantum models with nonclassical correlations may have better performance

as generative models on real-world data. The performance boost is also robust to practical training procedures and realistic performance metric considerations.

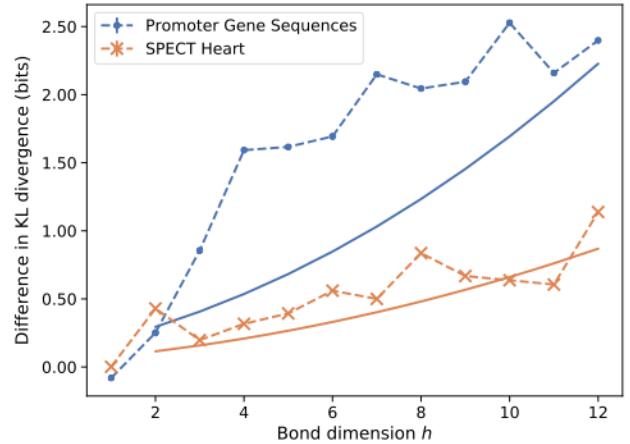


FIG. 8. Plotted with dashed lines are the improvements in KL divergence between the best-performing basis-enhanced 2-gram model and the best-performing classical hidden Markov model for each tested model size on the Promoter Gene Sequences and SPECT Heart data sets. The solid lines show the improvement needed to reject the null hypothesis in a likelihood-ratio test with 5σ confidence (see Sec. V C). Wherever the dashed lines are above their corresponding solid lines, the null hypothesis was rejected with more than 5σ confidence.

VI. CONCLUSION AND OUTLOOK

In this work, we have presented unconditional proof of the separation in expressive power between Bayesian networks and their minimal extension, basis-enhanced Bayesian quantum circuits. We showed that the origin of this separation is associated with quantum nonlocality and contextuality. Focusing on sequential models, we constructed examples via quantum nonlocality of a linear separation in k between k -gram models and their basis-enhanced version, and through quantum contextuality, a quasipolynomial separation in bond dimension for the hidden Markov model and its basis-enhanced version. In addition, we numerically tested this separation on standard data sets, showing that this separation holds even on practical data sets.

Although we focused on Bayesian networks, our approach can also be applied to more general models. Contextuality provides a general framework since the error model in Eq. (1) is independent of the normalization of probability distributions; therefore, our techniques can be applied to graphical models without well-defined transition probabilities along some edges of the graph. For example, Theorem 2 also works for deep Boltzmann machines, which is a much harder model than Bayesian networks in terms of computational cost. However, there is an intrinsic difficulty in extending Theorem 2 to get a separation with some non-energy-based neural networks (e.g., CNN, RNN [3]). The reason for this is that one hidden neuron in such kinds of models can take values over real numbers; thus, it could potentially carry infinite information, and our counting methods used in the proof of Theorem 2 do not directly apply. The possible extensions and applications of our approach to such models deserve further theoretical investigations. However, these models have to be implemented with finite precision in practice. In this case, our model still has a separation compared with these neural network models in terms of the amount of memory required when transmitting information between successive layers.

Our results establish a powerful connection between quantum foundations and machine learning. Since many traditional machine-learning models are based on the understanding and intuition from classical physics, they can be naturally characterized by noncontextual ontological models. Our study shows that quantum correlation can be a resource to enhance the expressive efficiency of these models even if the task is purely classical (e.g., by noting the similarity between contextuality in natural languages and quantum contextuality; see also Ref. [84] for the connection between natural language and quantum-inspired models from a different point of view). Our work opens new avenues for using ideas from quantum foundations to develop novel machine-learning models based on MPS (and its equivalence, OOMs, as we mentioned near the end of Sec. II), treelike tensor networks, or the multiscale entanglement renormalization ansatz (MERA)

[31–34,85,86]. Because of the form of our error model, the separation between classical and quantum-inspired models is mainly due to the non-negative versus complex numbers instead of the normalization in obtaining the probability. As an example, the general MPS has the same separation in terms of bond dimension compared with the stochastic MPS [87], which is a kind of MPS with non-negative numbers. In addition, we expect that the concepts of quantum correlations can be used to provide theoretical foundations for other quantum-inspired classical models and quantum machine-learning models.

Finally, our work provides new insights into designing practical quantum machine-learning algorithms that exhibit a quantum advantage in tackling machine-learning tasks; this can be achieved by starting from a successful classical machine-learning model and enhancing it with quantum correlations. This is a markedly different approach from most proposed quantum machine-learning models, in which generic quantum circuits are considered. Such structure is important since unstructured circuits tend to encounter challenges in training, including encountering barren plateaus in the loss landscape [51–53] and finding poor local minima [56]. Although the examples we used here can be efficiently simulated classically, some of them require the use of quantum machines during the training stage. Furthermore, the example models we used in this work are subclasses of more general, sequential, quantum generative models, involving quantum circuits with sequential (adaptive or nonadaptive) measurements, which are not classically simulable. It can be expected that the ideas of contextuality presented here can be extended to these cases to achieve a quantum advantage beyond classical simulability.

ACKNOWLEDGMENTS

We thank Stephen Bartlett, Dongling Deng, Weikang Li, Zhengfeng Ji, Jordi Tura, and Seth Lloyd for helpful discussions. X. G. is supported by the Postdoctoral Fellowship in Quantum Science of the Harvard-MPQ Center for Quantum Optics, the Templeton Religion Trust Grant No. TRT 0159, and by the Army Research Office under Grant No. W911NF1910302 and MURI Grant No. W911NF2010082. E. R. A. is supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. 4000063445, and a Lester Wolfe Fellowship and the Henry W. Kendall Fellowship Fund from MIT. S.-T. W. is partially supported by Air Force STTR under Grant No. FA8750-20-P-1708. J. I. C. acknowledges funding from ERC Advanced Grant QUENOCOBA under the EU Horizon 2020 program (Grant Agreement No. 742102), and the D-A-CH Lead-Agency Agreement through Project No. 414325145(BEYOND C). M. D. L. acknowledges funding from NSF CUA Grant No. PHY-1734011, NSF Grant No. PHY-2012023, ARO MURI Grant No. W911NF2010082, and DARPA ONISQ Grant No. W911NF2010021.

APPENDIX A: RELATIONS AMONG VARIOUS MACHINE-LEARNING MODELS

Deep belief nets have the form

$$p(v, h_1, h_2, \dots, h_d) = p(v|h_1) \dots p(h_{d-2}|h_{d-1})p(h_{d-1}, h_d); \quad (\text{A1})$$

this is exactly the form of a 2-gram model if the hidden variables are also observed.

A k -gram model can be simulated by a HMM with L^{k-1} hidden variables per site, where L is the vocabulary length of the k -gram model; this can be done straightforwardly by combining sets of $k-1$ sites in the k -gram model into one site in the HMM. The L^{k-1} possible values of these sites in the k -gram model map to each of the L^{k-1} hidden variables in the HMM.

APPENDIX B: DETAILS OF BASIS-ENHANCED BAYESIAN QUANTUM CIRCUITS

1. Mapping between Bayesian networks and quantum circuits

In the following, we give the explicit construction of the mapping between Bayesian networks and quantum circuits (see Fig. 9 as an illustration):

(i) Bayesian networks \Rightarrow BQCs. Each node corresponds to a qubit. According to the direction of

edges in the graph, assign an order for these qubits. Then, do the following steps in order. If the node has no parent, prepare the corresponding qubit as $|\psi_i\rangle$ such that $|\langle x_i|\psi_i\rangle|^2 = p(x_i)$. Otherwise, prepare the corresponding qubit as $|\psi_i\rangle$ and apply U (parents of x_i) on the corresponding qubits such that

$$|\langle x_i|U(\text{parents of } x_i)|\psi_i\rangle|^2 = p(x_i|\text{parents of } x_i). \quad (\text{B1})$$

Notice that there are no other operations between U (parents of x_i) and $|\psi_i\rangle$ because of the following: (i) If there were a target unit V in another uniformly controlled gate, U and V could be merged into a single uniformly controlled gate, and (ii) the order guarantees there is no control unit before a target unit.

(ii) BQCs \Rightarrow Bayesian networks. First, we assign a directed acyclic graph to the BQC. Each qubit corresponds to a node, and we draw an arrow from node x_i to node x_j if and only if there exists a uniformly controlled gate with a control unit on qubit i and a target unit on qubit j . Then, we assign each node transition probabilities in the following way: If there is no target unit on qubit i , assign $p(x_i) = |\langle x_i|\psi_i\rangle|^2$ for the corresponding node; otherwise, assign a transition probability for this node according to Eq. (B1).

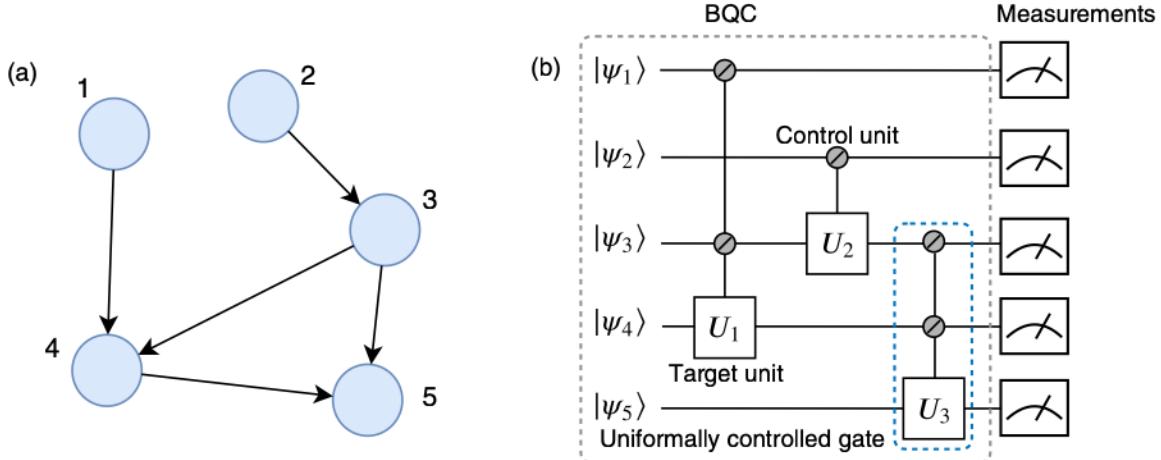


FIG. 9. Bayesian network on a generic graph and its associated Bayesian quantum circuit. (a) Example of a Bayesian network on a generic graph. Each node x_i corresponds to a transition probability $p(x_i|\text{parents of } x_i)$; e.g., the node x_4 corresponds to $p(x_4|x_1, x_3)$. If there is no parent for the node x_i , it simply corresponds to a marginal probability $p(x_i)$; e.g., the node x_2 corresponds to $p(x_2)$. The joint probability distribution defined by the Bayesian network is the product of the transition or marginal probability over all of the nodes. For this example, the probability distribution is $p(x_1, x_2, x_3, x_4, x_5) = p(x_1)p(x_2)p(x_3|x_2)p(x_4|x_1, x_3)p(x_5|x_3, x_4)$. (b) Associated Bayesian quantum circuit. Each node in the Bayesian network corresponds to a qubit, and each uniformly controlled gate corresponds to a transition probability. For example, the uniformly control- U_2 gate corresponds to $p(x_4|x_1, x_3)$. The marginal probabilities $p(x_1), p(x_2)$ have been absorbed into $|\psi_1\rangle, |\psi_2\rangle$. Because all Bayesian networks are associated with directed acyclic graphs, an ancestor of a node cannot also be a child of the same node. This means, in Bayesian quantum circuits, there cannot be a target unit after a control unit. Furthermore, the target unit can only involve one qubit in a uniformly controlled gate. Measuring the output qubits of Bayesian quantum circuits in the computational basis will produce the same probability distribution as the corresponding Bayesian network.

2. Efficient implementation using multiqubit collective gates

The implementation of uniformly controlled gates is not efficient, in general [67]. This is true even if we have the ability to implement collective gates that are native, for instance, to Rydberg-based quantum platforms [e.g., implementing quantum fan-out gates with k control units [88] as shown in Fig. 10(a)]. Even though these gates are very powerful [89], it is unclear how to implement general uniformly controlled gates more efficiently (in terms of scaling with k). However, in almost all machine-learning models, the transition probabilities have specific forms when k is large. In particular, transition probabilities usually take the form

$$p(x_{k+1}|x_1, \dots, x_k) = f\left(x_{k+1} \left(\sum_{i=1}^k w_i x_i \right)\right); \quad (\text{B2})$$

that is, the dependence on the parent variables is linear in a nonlinear function f , in general. For example, in deep belief nets, $f(y) = e^{-\beta y} / (1 + e^{-\beta y})$. Here, we give a construction showing how to implement such uniformly controlled gates approximately such that the number of elementary collective gates does not depend on k .

In the following, we show how to implement the transition probability shown in Eq. (B2) in BQCs with a circuit depth independent of k , using collective gates. We define $\theta = \sum w_i x_i$, such that according to Eq. (B2), $p(0|x_1, \dots, x_k) = f(0)$ and $p(1|x_1, \dots, x_k) = f(\theta)$. Thus, we have a normalization condition $f(0) + f(\theta) = 1$. We introduce the notation $\langle \cdot \rangle_m$ as a binary representation of \cdot up to the m th digit. We also introduce $\tilde{\theta}$ as an approximation of θ , with binary representation $\langle \theta \rangle_{d_1}$. We then use the following procedure to implement the transition:

$$\begin{aligned} |0\rangle^{\otimes d_1} |0\rangle^{\otimes d_2} |0\rangle &\rightarrow |\langle \theta \rangle_{d_1} \rangle |0\rangle^{\otimes d_2} |0\rangle \quad [\text{phase estimation algorithm, } O(d_1^2) \text{ gates}] \\ &\rightarrow |\langle \theta \rangle_{d_1} \rangle |\arcsin \sqrt{f(\tilde{\theta})} \rangle_{d_2} |0\rangle \quad [\text{classical computing, usually } O(\text{poly}(d_2^2)) \text{ gates}] \\ &\rightarrow |\langle \theta \rangle_{d_1} \rangle |\arcsin \sqrt{f(\tilde{\theta})} \rangle_{d_2} (\sqrt{f(0)}|0\rangle + \sqrt{f(\theta)}|1\rangle) \quad (\text{controlled rotation along } x \text{ axis}) \\ &\rightarrow |0\rangle^{\otimes d_1} |0\rangle^{\otimes d_2} (\sqrt{f(0)}|0\rangle + \sqrt{f(\theta)}|1\rangle) \quad (\text{uncomputing}). \end{aligned} \quad (\text{B3})$$

The precision of the transition is determined by d_1 and d_2 . Note that d_1 determines the precision of the input of the function f , and d_2 determines the effect of truncation for the function $\arcsin \sqrt{f(\cdot)}$. The total error is bounded by

$$\epsilon = \max_{\theta} |(\arcsin \sqrt{f(\theta)})'| 2^{-d_1} + 2^{-d_2}. \quad (\text{B4})$$

Therefore, if the derivative is bounded by a constant [for example, in the case of $f(\theta) = e^{-\beta\theta} / (1 + e^{-\beta\theta})$, the derivative is bounded by β], d_1 and d_2 can be taken to be

$O(\log(1/\epsilon))$ such that the depth is bounded by $\text{poly}(\log(1/\epsilon))$, which is independent of k and only depends on the precision ϵ .

3. Exponential separation of expressive power between BBQCs and Bayesian networks based on computational complexity theory

The proof of the exponential expressive power of BBQCs is a slight modification of the proof for quantum generative models (QGMs) detailed in Ref. [16].

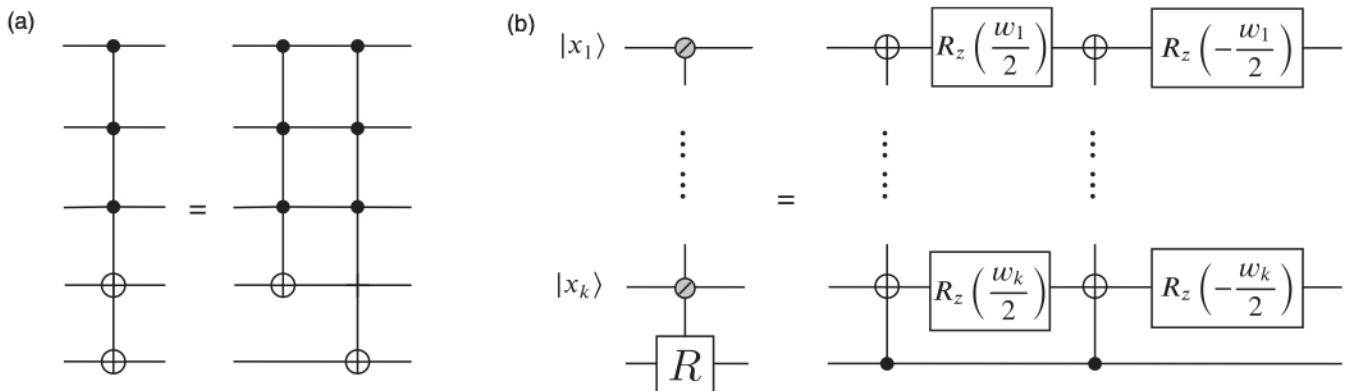


FIG. 10. Basic elements in Bayesian quantum circuits implemented by multiqubit collective gates. (a) “Basic” collective gate. (b) Implementation of control- $R = R_z(\sum_i w_i x_i)$ gate. This gate can be implemented directly on a Rydberg-atom-based platform.

Theorem 3. (Ref. [16]) There exists a BBQC with n qubits such that, if any Bayesian networks with a polynomial number of parameters in n could approximate it under the multiplicative error, the polynomial hierarchy in computational complexity theory would collapse.

For completeness, we give a brief review of the proof. First, we give a brief introduction of related concepts. Second, we introduce a specific BBQC, which is used to separate the expressive power between the classical and quantum models. Third, we give a sketch of the proof. See Ref. [16] for more details.

a. Related computational complexity classes

The polynomial hierarchy is a hierarchy of complexity classes that generalize \mathbf{P} and \mathbf{NP} , and are denoted as $\Sigma_0^p, \Sigma_1^p, \Sigma_2^p, \dots$. Here, $\Sigma_0^p = \mathbf{P}$, $\Sigma_1^p = \mathbf{NP}$, and $\Sigma_{i+1}^p = \mathbf{NP}^{\Sigma_i^p}$, where $\mathbf{NP}^{\Sigma_i^p}$ is called \mathbf{NP} relative to Σ_i^p . \mathbf{NP} denotes problems that can be verified in polynomial time by a Turing machine, and $\mathbf{NP}^{\Sigma_i^p}$ denotes problems that can be verified in polynomial time by a Turing machine that is equipped with an oracle that can solve any Σ_i^p problems in one step. A detailed discussion can be found in Ref. [90] or in the recent review article on quantum supremacy [91]. It is widely believed that the polynomial hierarchy does not collapse, which means $\Sigma_i^p \neq \Sigma_{i+1}^p$ (which implies $\Sigma_i^p \neq \Sigma_{i+j}^p$ for any constant $j > 0$).

b. Basis-enhanced Bayesian network used in the proof

Here, we give a construction of a basis-enhanced Bayesian network such that approximately computing the probability of a specific configuration up to the multiplicative error is $\#\mathbf{P}$ -hard.

This BBQC begins as a cluster state on a square lattice. The corresponding Bayesian network is drawn as the graph shown in Fig. 11. Then, we use the measurement basis

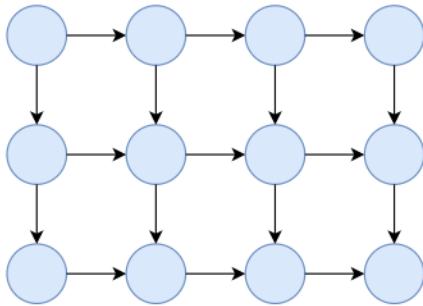


FIG. 11. A cluster state in a Bayesian quantum circuit. A cluster is generated from initial state $|+\rangle^{\otimes n}$ by applying control-Z gates between each pair of neighbors on a square lattice. We may assign each edge an arrow in the way shown in this figure in order to get a directed acyclic graph. We can see that this circuit is a Bayesian quantum circuit by checking Definition 1 and noticing that, in a control-Z gate, there is no need to distinguish between control and target units.

shown in Ref. [92]. One of the important properties of this construction is its “single-instance hardness,” which means there is only one measurement basis for any fixed size; i.e., the probability distribution $q(\mathbf{x})$ only depends on the size of the lattice. We demand this property because in the proof of exponential expressive power, we associate a probability distribution with a problem consisting of non-negative numbers as outputs: \mathbf{x} specifies an instance of the problem, and the task is to compute the probability given a specific \mathbf{x} , i.e., $q(\mathbf{x})$, to the multiplicative error. Thus, the complexity of a probability distribution is defined as the complexity of the associated problem.

The proof of exponential expressive power works for any efficiently computable classical model. Thus, it also works for any neural networks.

c. Sketch of the proof

The key to separating the complexity of the classical and quantum models is formalizing a sign problem caused by quantum interference: Approximately computing (up to the multiplicative error) a summation of many non-negative numbers is easier than the summation of many complex or real numbers. This can be done via Stockmeyer’s theorem [93] (see Ref. [16] for an introduction oriented to the proof here); the former is inside Σ_2^p , and the latter is $\#\mathbf{P}$ -hard. The same reasoning has been used to separate QGMs and the general probabilistic graphical model [16]. Here, we only give a sketch of the proof.

Assume there exists a Bayesian network that generates the joint probability $p(\mathbf{x}, \mathbf{y}, \mathbf{z})$ such that the conditional probability $\sum_{\mathbf{y}} p(\mathbf{x}, \mathbf{y} | \mathbf{z})$ approximates $q(\mathbf{x})$ to the multiplicative error. We can use Stockmeyer’s theorem to prove that, based on this assumption and supposing that the parameters of the network are given, approximating $q(\mathbf{x})$ to the multiplicative error is in Σ_2^p . We should keep in mind that the probability defines a problem, with \mathbf{x} specifying an instance of the problem.

However, though we can show that it is possible to approximate q to the multiplicative error, the proof is not constructive. More concretely, “/poly” denotes that, for any fixed input size (the length of \mathbf{x}), there exists a polynomial-sized classical circuit that computes all the instances of the problem, but the circuit may not be efficiently constructed [90, 94]. In Appendix B 3 b, we construct a BBQC such that computing $q(\mathbf{x})$ to the multiplicative error is $\#\mathbf{P}$ -hard. Thus, assuming the efficient representation of the BBQC via classical Bayesian networks, we (roughly) obtain

$$\#\mathbf{P} \subseteq \Sigma_2^p / \text{poly}. \quad (B5)$$

This implies that the polynomial hierarchy would collapse to the third level, as more formally shown in Ref. [16] (which follows from a modification of the reasoning of the proof of Theorem 3 in Ref. [95]).

4. Algorithms for inference and learning

There are mainly two computational problem associated with a generative model. One is inference, i.e., how to extract useful information from the representation of the generative model. Making inference on a generative model usually means computing marginal probabilities or conditional probabilities, or performing maximum likelihood estimation. With this, we can make predictions for new data after getting an approximately correct representation of a data distribution p_D . Later, we will give examples to show the applications of computing conditional probabilities.

The other computational problem is training (or learning), i.e., how to determine parameters of the generative model from training data in order to approximate p_D . Training usually means minimizing the KL divergence

$$D(p_D||p_\theta) = \sum_v p_D(v) \log \left(\frac{p_D(v)}{p_\theta(v)} \right) \quad (\text{B6})$$

between p_D and p_θ , the distribution of the generative model, with the whole parameter set denoted by $\theta = (\theta_1, \dots, \theta_{\text{poly}(n)})$. The θ -dependent part of $D(p_D||p_\theta)$ can be expressed as

$$D(\theta) \equiv -\frac{1}{N} \sum_{v \in \text{training data set}} \log p_\theta(v), \quad (\text{B7})$$

where N denotes the total number of data [$1/N$ approximates $p_D(v)$] and the summation is over all the training data or a batch of data for stochastic optimization. As the number of parameters is bounded by $\text{poly}(n)$, the required data size N is typically bounded by $\text{poly}(n)$ [1]. We can also understand optimizing $D(\theta)$ as maximum likelihood estimation since $D(\theta)$ is proportional to the log-likelihood $\log(\prod_v q_\theta(v))$. It is worth mentioning that, in addition to $D(\theta)$ with $v = (x, y)$, it is also usual to adopt the following loss function for supervised learning [96],

$$L(\theta) \equiv -\sum_{v \in \text{training data set}} \log q_\theta(y|x), \quad (\text{B8})$$

since it is the log-likelihood $\log(\prod_v q_\theta(y|x))$. Typically, we minimize these loss functions via a so-called optimizer, usually the gradient descent method [3] with a proper learning rate (the step length for updating parameters) or its variations, like adding a stochastic term, adjusting the learning rate adaptively, utilizing training data in batches, and so on.

a. Heuristic quantum algorithms for inference and learning

Even for classical Bayesian networks, the training and inference problems are computationally hard for quantum computers [79,80]. However, there are a number of

proposed heuristic and approximate algorithms that work well in practice in tackling these computation problems for classical algorithms [81].

BBQCs have a similar problem in that exact training and inference are computationally difficult, and it is natural to propose heuristic quantum algorithms. Since Bayesian networks are a special case of probabilistic graphical models [81], the quantum algorithm for the learning and inference problems in extensions of probabilistic graphical models [16] also works here. The idea is to convert the learning and inference problems to preparing ground states of a local Hamiltonian. The runtime of the quantum algorithm is proportional to the inverse of the energy gap, although we cannot guarantee that the energy gap scales as $1/\text{poly}(n)$. Other heuristic quantum algorithms more specific to Bayesian networks may also exist as in the classical case.

APPENDIX C: RELATIONS AMONG VARIOUS ERROR MODELS

For completeness, we also define two other error models. One is the multiplicative error:

$$|p(x) - q(x)| \leq \gamma q(x), \quad \forall x, \quad (\text{C1})$$

with γ being a constant smaller than $1/2$. Note that p approximating q under this error implies that $D_{\text{KL}}(p||q)$ is bounded by γ . Thus, p approximating q under this error model is a stronger requirement than a small KL divergence of p from q . This error model is used for our complexity-theory-based proof of a quantum advantage on general graphs in Appendix B 3.

For translation problems, the generative models usually only define conditional probabilities $p(y|x)$ for the classical model [e.g., the second model in Fig. 1(b)] and $q(y|x)$ for the quantum extension. The prior probability for x is unspecified, and we denote it as $p(x)$ and $q(x) = p(x)$. Then, the KL divergence is

$$\begin{aligned} D_{\text{KL}}(p||q) &= \sum_{x,y} p(x,y) \log \frac{p(x,y)}{q(x,y)} \\ &= \sum_x p(x) \left(\sum_y p(y|x) \log \frac{p(y|x)}{q(y|x)} \right). \end{aligned} \quad (\text{C2})$$

In order to avoid any assumptions on $p(x)$, we require that the quantity inside the brackets be bounded for any x . This implies that

$$q(y|x) = 0 \Leftrightarrow p(y|x) = 0, \quad \forall x, y. \quad (\text{C3})$$

Now, let us show that the multiplicative error is bounded, which implies that $D_{\text{KL}}(p||q)$ and $D_{\text{KL}}(q||p)$ are bounded; this, in turn, implies the error model in Eq. (1). We see that

$$\begin{aligned}
\left| \sum_x p(\mathbf{x}) \log \frac{p(\mathbf{x})}{q(\mathbf{x})} \right| &\leq \sum_x \left| p(\mathbf{x}) \log \left(1 + \frac{p(\mathbf{x})}{q(\mathbf{x})} - 1 \right) \right| \\
&\leq \sum_x \left| p(\mathbf{x}) \left(\frac{p(\mathbf{x})}{q(\mathbf{x})} - 1 \right) \right| \\
&\leq \sum_x p(\mathbf{x}) \gamma \\
&= \gamma.
\end{aligned} \tag{C4}$$

The second inequality comes from

$$|p(\mathbf{x}) - q(\mathbf{x})| \leq \gamma q(\mathbf{x}) \tag{C5}$$

with small γ . A similar proof holds for $D_{\text{KL}}(q||p)$ by exchanging p and q .

According to the definition of $D_{\text{KL}}(p||q)$, $q(\mathbf{x}) = 0$ implies that $p(\mathbf{x}) = 0$ in order to make $D_{\text{KL}}(p||q)$ bounded. According to the definition of $D_{\text{KL}}(q||p)$, $p(\mathbf{x}) = 0$ implies that $q(\mathbf{x}) = 0$ in order to make $D_{\text{KL}}(q||p)$ bounded. Thus, both $D_{\text{KL}}(p||q)$ and $D_{\text{KL}}(q||p)$ are bounded, which implies the error model in Eq. (1).

APPENDIX D: LEMMA PROOFS FOR THE MERMIN-PERES MAGIC SQUARE

First, we prove the following lemma:

Lemma 2. For any subset of stabilizer states s such that $|s| > 2^{n^2/4+7n/2}$, there exist three states such that some of their stabilizers form a Mermin-Peres magic square as shown in Table II.

Proof.—We write $|s| = a(n)b(n)c(n) + 1$, where $a(n) = 2^{n^2/4+3n/2}$, $b(n) = 4^n$, and $c(n) = 2$; their meaning will be clear later. Given $|s|$ stabilizer states, we can always transform these states to another set of stabilizer states by a Clifford circuit such that one of the states will become $|0\rangle^{\otimes n}$ and the other $a(n)b(n)c(n)$ states have the following form (see Ref. [97]):

$$|\psi\rangle \propto \sum_{x \in A} (-1)^{q(x)} i^{l(x)} |x\rangle, \tag{D1}$$

where A is an affine subspace of \mathbb{Z}_2^n and $q(x)$ and $l(x)$ are quadratic and linear functions on \mathbb{Z}_2 and \mathbb{Z}_4 , respectively. The state is determined by A , q , l .

We denote $a(n)$ as the number of different A : An affine subspace is composed of a linear subspace and a displacement, and thus there are at most

$$a(n) \leq \sum_k 2^{(n-k+1)k} \times 2^n = 2^{n^2/4+3n/2} \tag{D2}$$

TABLE II. Mermin-Peres magic square.

Computational states	$(-1)^f I_1 Z_2 Z^{a_3+e_3}$	$(-1)^g Z_1 I_2 Z^{b_3+d_3}$	$(-1)^{f+g} Z_1 Z_2 Z^{a_3+b_3+e_3+d_3}$
First graph state	$X_1 I_2 Z^{a_3}$	$I_1 X_2 Z^{b_3}$	$X_1 X_2 Z^{a_3+b_3}$
Second graph state	$X_1 Z_2 Z^{e_3}$	$Z_1 X_2 Z^{d_3}$	$-X_1 X_2 Z_1 Z_2 Z^{e_3+d_3}$

possible A , where the first term involving summation over k is the number of linear subspaces (where k is the dimension of the subspace—see Theorem 2.14 of Ref. [98]) and the second term 2^n is the number of possible displacements.

According to the pigeonhole principle, we can prove that we now have the $|0\rangle^{\otimes n}$ state and at least $b(n)c(n)$ states belonging to the same affine subspace A . Those $b(n)c(n)$ states only differ by q and l . Using C_{CNOT} and C_x , which are circuits only composed of CNOT and Pauli x gates, respectively, these states can be transformed to be of the form

$$\sum_{u \in \{0,1\}^k} (-1)^{\bar{q}(u)} i^{\bar{l}(u)} |u\rangle |0\rangle^{\otimes(n-k)}, \tag{D3}$$

which are graph states over the first k qubits. These circuits simultaneously transform $|0\rangle^{\otimes n}$ to a state of the form $|z_1 z_2 \dots z_n\rangle$.

Denote $b(n)$ as the number of \bar{l} which is no greater than 4^n (where the worst case is $k = n$). Then, we have the $|z_1 z_2 \dots z_n\rangle$ state and at least $c(n)$ graph states after applying S or Z gates to eliminate \bar{l} . As long as $c(n) \geq 2$, we can always find two graph states such that there exists a pair of vertices where there is no edge for the first graph and there is an edge for the second graph. Without loss of generality, we may assume this pair is comprised of qubits 1 and 2.

The first graph state (without an edge between qubits 1 and 2) has stabilizer generators $X_1 I_2 Z^{a_3}$ and $I_1 X_2 Z^{b_3}$, and the second has generators $X_1 Z_2 Z^{c_3}$ and $Z_1 X_2 Z^{d_3}$, where a_3 , b_3 , c_3 , d_3 are $n - 2$ -dimensional vectors on \mathbb{Z}_2 . The computational state has generators $\langle (-1)^{e_i} Z_i \rangle$, so it could have any Pauli Z -type stabilizers up to \pm signs. Then, we have the Mermin-Peres magic square given by Table II. In this table, operators in each row commute with each other since they are chosen from stabilizers of the same quantum states. It is also easy to check that operators in each column commute with each other. The first two Pauli stabilizers in each observable form the “traditional” Mermin square, and the Pauli Z stabilizers after the first two qubits do not change the commutation relations between observables. Thus, this table forms a Mermin-Peres magic square and thus exhibits contextuality. ■

Second, we prove the following lemma:

Lemma 3. If three stabilizer states and a subset of their stabilizers form a Mermin-Peres magic square as shown in Table I, the intersection of their support in an ontological theory should be empty in order to be consistent with quantum mechanics.

Proof.—The proof basically follows the discussion in the main text with the example in Eq. (11) but written in a more general way. Assume there is a common λ in the intersection among the supports of the three states $|\psi_1\rangle$, $|\psi_2\rangle$, and $|\psi_3\rangle$.

It is simple to show the following two equations:

$$\begin{aligned} \frac{1 - ABab}{2} |\psi_3\rangle &= 0, \quad \text{or equivalently,} \\ \frac{1 + ABab}{2} |\psi_3\rangle &= |\psi_3\rangle, \end{aligned} \quad (\text{D4})$$

and

$$\begin{aligned} \frac{1 + ABab}{2} |\psi_1\rangle \perp \frac{1 + ABab}{2} |\psi_2\rangle, \quad \text{or equivalently,} \\ \langle \psi_1 | \frac{1 + ABab}{2} |\psi_2\rangle &= \langle \psi_1 | \frac{1 + Aa}{2} \frac{1 + ABab}{2} \frac{1 + Bb}{2} |\psi_2\rangle \\ &= 0. \end{aligned} \quad (\text{D5})$$

The first set of equalities forces the measurement result of λ to deterministically be +1. The second set of equalities shows that the resulting states of $|\psi_1\rangle$ and $|\psi_2\rangle$ after measuring $ABab$ and getting +1 are orthogonal. Thus, there is a contradiction. ■

APPENDIX E: ROBUST SEPARATION OF k -GRAM MODEL UNDER l_1 DISTANCE

Here, we prove that any k -gram model with the probability distribution p with $k < n/6$ cannot approximate a particular basis-enhanced 2-gram model with the probability distribution q to l_1 distance smaller than 1/288, i.e.,

$$\sum_x |p(x) - q(x)| \geq 1/288. \quad (\text{E1})$$

For simplicity, we assume $n = 3 + 4k$. The key to proving the separation between the quantum extension and its classical counterpart is through a Bell test of the GHZ state through measurements in the X and Y bases. By measuring the remaining qubits, we obtain a GHZ state up to three single-qubit Clifford gates. However, as we restricted measurement to the X and Y bases, this does not always hold. The following lemma gives the probability of still having nonlocality.

Lemma 4. The probability of measuring the remaining qubits to get a GHZ state up to Pauli and S gates is larger than 1/9. We call this measurement a *GHZ-type measurement*.

Proof.—Suppose the measurement basis and results for the remaining qubits are $b_1 s_1 b_2 s_2 \dots$ with equal probability. Then, the resulting state is

$$\sigma_1 C_1 \otimes \sigma_3 \otimes \sigma_2 C_2 |\text{GHZ}\rangle, \quad (\text{E2})$$

where σ_i is a Pauli matrix and

$$\begin{aligned} C_1 &= HS^{s_1} H \cdot S^{s_2} \cdots \cdot HS^{s_{2k-1}} H \cdot S^{s_{2k}}, \\ C_2 &= HS^{s_{2k+1}} H \cdot S^{s_{2k+2}} \cdots \cdot HS^{s_{4k-1}} H \cdot S^{s_{4k}}. \end{aligned} \quad (\text{E3})$$

We only need to prove that the probability of C_1 equaling I or S up to a Pauli matrix is at least 1/3.

All of the single-qubit Clifford gates can be represented as permutations in S_3 among single-qubit Pauli matrices up to an unimportant phase factor. Note that HSH and S can be regarded as (12) and (23), which are generators of S_3 . Starting from I , each time we apply $HS^{s_{2i-1}} H \cdot S^{s_{2i}}$ with probability 1/4 for all of the choices of s_{2i-1}, s_{2i} , we obtain a random walk among the six group elements of S_3 . The transfer matrix is

$$\begin{pmatrix} \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} & 0 \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \\ \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & 0 & \frac{1}{4} \\ 0 & \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \\ 0 & \frac{1}{4} & 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{4} \end{pmatrix}. \quad (\text{E4})$$

By solving for the eigenstates and eigenvalues, and choosing an initial state of $(1,0,0,0,0,0)$, we find that after k steps, the probability to get I and S up to Pauli operators [which are (1) and (12) in S_3] is given by

$$\frac{1}{3} + \frac{2}{3} 4^{-k} > \frac{1}{3}, \quad (\text{E5})$$

which proves the lemma. ■

Lemma 5. For distributions p and q , and any positive number c ,

$$\sum_x |p(x) - cq(x)| \geq \frac{\sum_x |p(x) - q(x)|}{2}. \quad (\text{E6})$$

Proof.—Denote $\delta = \sum_x |p(x) - q(x)|$. We consider the following two cases:

(i) $|1 - c| \geq \delta/2$:

$$\begin{aligned} \sum_x |p(x) - cq(x)| &\geq \left| \sum_x p(x) - cq(x) \right| \\ &= |1 - c| \\ &\geq \delta/2. \end{aligned} \quad (\text{E7})$$

(ii) $|1 - c| \leq \delta/2$:

$$\begin{aligned} \sum_x |p(x) - cq(x)| &\geq \sum_x |p(x) - q(x)| - |(1 - c)q(x)| \\ &= \delta - |1 - c| \\ &\geq \delta/2. \end{aligned} \quad (\text{E8})$$

Lemma 6. Denote the measurement bases and results for the three chosen qubits as $\mathbf{a} = (a_1, a_2, a_3)$ and $\mathbf{t} = (t_1, t_2, t_3)$, respectively. Then,

$$\sum_{\mathbf{a}, \mathbf{t} \in \text{sol}_a} q(\mathbf{a}) p(\mathbf{t}|\mathbf{a}, \text{GHZ-type}) \leq \frac{7}{8}, \quad (\text{E9})$$

where sol_a means \mathbf{t} satisfies Eq. (5) up to flips of some b_i and s_i determined by the GHZ-type measurement basis and results.

Proof.—The probability distribution in Eq. (6) could also be understood as follows: There is a probability distribution $p(\lambda)$, and each λ determines t_i , i.e., $t_i = t_i(s_i, \lambda)$. Because the GHZ test cannot be described by a local hidden variable theory, there exists at least one assignment of \mathbf{a} given λ such

that $\mathbf{t} = \mathbf{t}(s, \lambda) \notin \text{sol}_a$. Assuming this and $q(\mathbf{a}) = 1/8$, we have that

$$\begin{aligned} \sum_{\mathbf{a}, \mathbf{t} \in \text{sol}_a} q(\mathbf{a}) p(\mathbf{t}|\mathbf{a}) &= \frac{1}{8} \sum_{\mathbf{a}} \sum_{\mathbf{t} \in \text{sol}_a} p(\mathbf{t}|\mathbf{a}) \\ &= \frac{1}{8} \int_{\lambda} p(\lambda) \sum_{\mathbf{a}} \mathbf{1}(\mathbf{t}(\lambda, \mathbf{a}) \in \text{sol}) \\ &\leq \frac{7}{8}, \end{aligned} \quad (\text{E10})$$

where $\mathbf{1}(\cdot)$ is the indicator function, which is 1 if the condition \cdot holds and 0 otherwise. ■

Combining the above lemmas, we now show that Eq. (E1) holds. First,

$$\begin{aligned} \sum_{\mathbf{a}, \mathbf{b}, \mathbf{s}, \mathbf{t}} |p(\mathbf{a}, \mathbf{b}, \mathbf{s}, \mathbf{t}) - q(\mathbf{a}, \mathbf{b}, \mathbf{s}, \mathbf{t})| &= \sum_{\mathbf{a}, \mathbf{b}, \mathbf{s}, \mathbf{t}} |p(\mathbf{b}, \mathbf{s}) p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{b}, \mathbf{s}) q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})| \\ &= \left(\sum_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} + \sum_{\mathbf{b}, \mathbf{s} \notin \text{GHZ-type}} \right) \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{b}, \mathbf{s}) p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{b}, \mathbf{s}) q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})| \\ &\geq \sum_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{b}, \mathbf{s}) p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{b}, \mathbf{s}) q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})| \\ &\geq \sum_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} p(\mathbf{b}, \mathbf{s}) \sum_{\mathbf{a}, \mathbf{t}} \left| p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - \frac{q(\mathbf{b}, \mathbf{s})}{p(\mathbf{b}, \mathbf{s})} q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) \right| \\ &\geq \sum_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} \frac{p(\mathbf{b}, \mathbf{s})}{2} \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})| \\ &\geq \sum_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} \frac{p(\mathbf{b}, \mathbf{s})}{2} \min_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})| \\ &= \frac{1}{18} \min_{\mathbf{b}, \mathbf{s} \in \text{GHZ-type}} \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s}) - q(\mathbf{a}, \mathbf{t}|\mathbf{b}, \mathbf{s})|. \end{aligned} \quad (\text{E11})$$

Making the minimization over GHZ types implicit and noting that $q(\mathbf{t}|\mathbf{a}, \text{GHZ-type}) = 1$ for $\mathbf{t} \in \text{sol}_a$ finally yields

$$\begin{aligned} \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{a}, \mathbf{t}) - q(\mathbf{a}, \mathbf{t})| &= \sum_{\mathbf{a}, \mathbf{t}} |p(\mathbf{a}) p(\mathbf{t}|\mathbf{a}) - q(\mathbf{a}) q(\mathbf{t}|\mathbf{a})| \\ &= \sum_{\mathbf{a}} q(\mathbf{a}) \sum_{\mathbf{t}} |q(\mathbf{t}|\mathbf{a}) - \frac{p(\mathbf{a})}{q(\mathbf{a})} p(\mathbf{t}|\mathbf{a})| \\ &\geq \frac{1}{2} \sum_{\mathbf{a}} q(\mathbf{a}) \sum_{\mathbf{t}} |q(\mathbf{t}|\mathbf{a}) - p(\mathbf{t}|\mathbf{a})| \\ &= \frac{1}{2} \left(1 - \sum_{\mathbf{a}, \mathbf{t} \in \text{sol}_a} q(\mathbf{a}) p(\mathbf{t}|\mathbf{a}) \right) \\ &\geq \frac{1}{16}. \end{aligned} \quad (\text{E12})$$

With the last inequality, we arrive at the separation of 1/288 under the l_1 distance.

APPENDIX F: HHM TO SIMULATE FIG. 2(a)

The process described in Fig. 2 could also be understood by the HMM model described in Sec. IV C. The total number of possible quantum states involved is an upper bound for the number of hidden variables in HMM. The state stored in the second qubit of each pair in Fig. 2, with the evolution driven by measuring the first qubit in each

pair, only involves single-qubit stabilizer states. The total number is thus $6(|0\rangle, |1\rangle, |+\rangle, |-\rangle, |+_i\rangle, |-_i\rangle)$.

APPENDIX G: SUPPLEMENTAL NUMERICS FIGURES

In Fig. 12, we plot the tensor network representation of the derivative of the loss function Eq. (20). Furthermore,

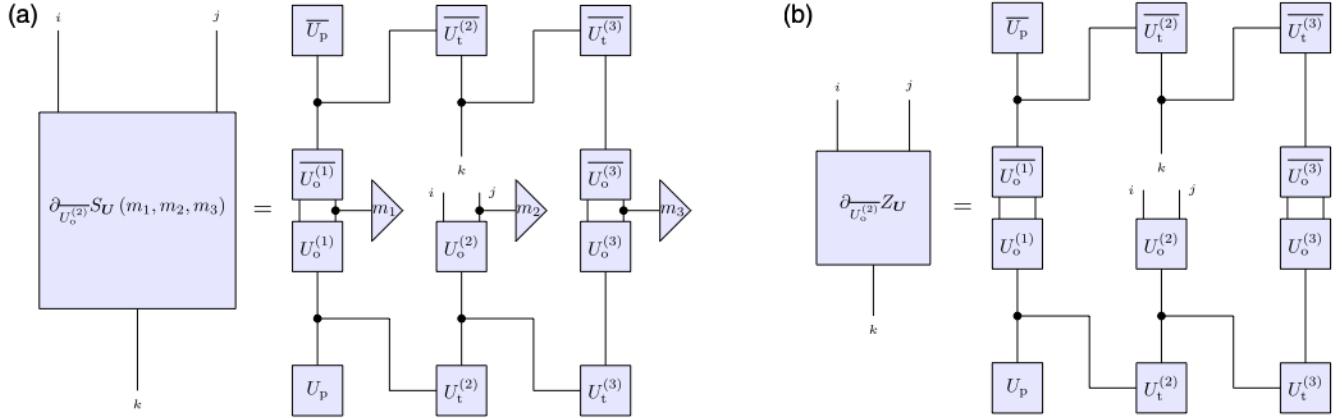


FIG. 12. (a) Derivative of $S_U(m)$ with respect to $\overline{U_o^{(2)}}$, using the network given in Fig. 6(b) as an illustrative example. (b) Derivative of Z_U with respect to $\overline{U_o^{(2)}}$, using the network given in Fig. 6(b) as an illustrative example.

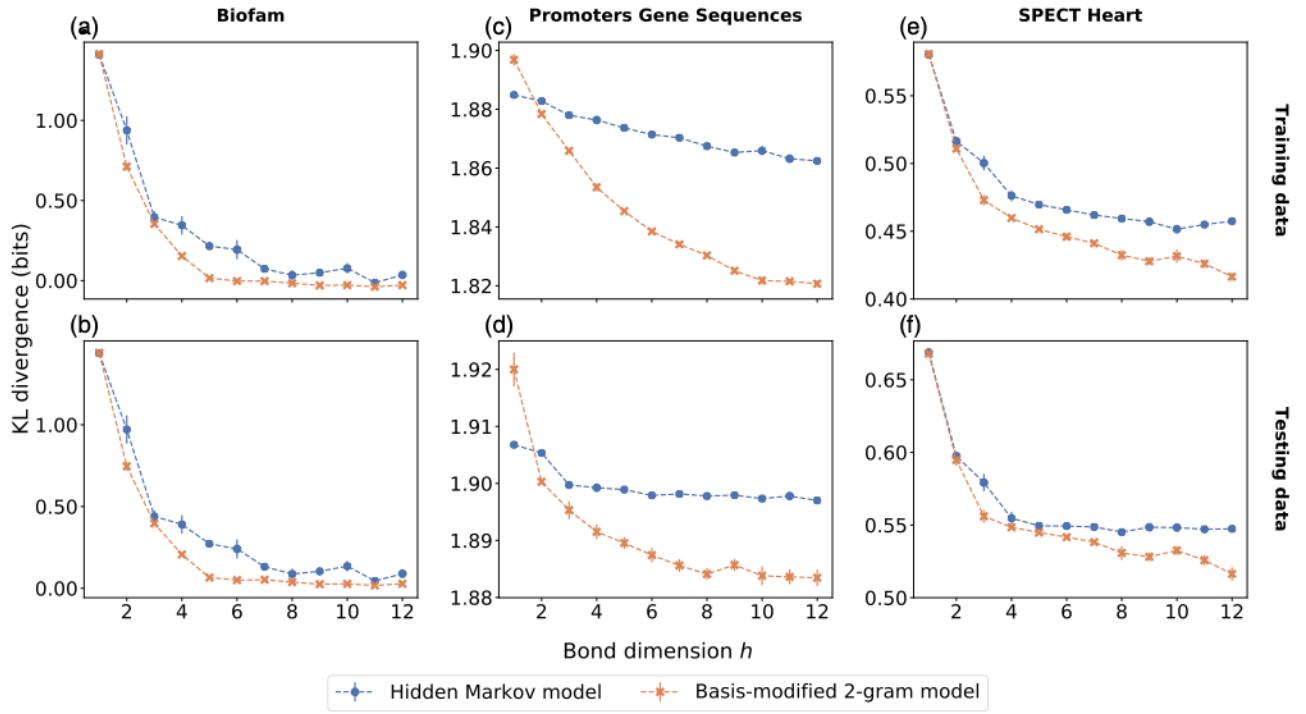


FIG. 13. Average performances over ten trials of the classical HMM (blue circles) and the basis-enhanced 2-gram model (orange crosses) on (a,b) the biofam, (c,d) Promoter Gene Sequences, and (e,f) SPECT Heart data sets. The first row plots the performance on the training data, and the second the performance on the withheld testing data. The basis-enhanced 2-gram model, on average, performed better than the classical model even on the simple biofam data set, implying more consistent performance in the basis-enhanced 2-gram model. Error bars denote 1 standard error of the mean over ten trials. Dashed lines are to aid the eye.

in Fig. 13, we plot the average performance of both the trained HMM and the basis-enhanced 2-gram model. Note that the performance separation between quantum and classical models persists even when considering the average performance over many runs. A slight separation is observed even for the biofam data set, implying more consistent performance in the basis-enhanced 2-gram model.

- [1] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning: From Theory to Algorithms* (Cambridge University Press, Cambridge, England, 2014).
- [2] C. M. Bishop, *Pattern Recognition and Machine Learning* (Springer, New York, 2006).
- [3] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [4] F. Arute *et al.*, *Quantum Supremacy Using a Programmable Superconducting Processor*, *Nature (London)* **574**, 505 (2019).
- [5] H.-S. Zhong *et al.*, *Quantum Computational Advantage Using Photons*, *Science* **370**, 1460 (2020).
- [6] Y. Wu *et al.*, *Strong Quantum Computational Advantage Using a Superconducting Quantum Processor*, *Phys. Rev. Lett.* **127**, 180501 (2021).
- [7] Q. Zhu *et al.*, *Quantum Computational Advantage via 60-Qubit 24-Cycle Random Circuit Sampling*, *Science bulletin* **67**, 280 (2021).
- [8] M. Schuld, I. Sinayskiy, and F. Petruccione, *An Introduction to Quantum Machine Learning*, *Contemp. Phys.* **56**, 172 (2015).
- [9] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Quantum Machine Learning*, *Nature (London)* **549**, 195 (2017).
- [10] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchytskyy, and R. Melko, *Quantum Boltzmann Machine*, *Phys. Rev. X* **8**, 021050 (2018).
- [11] 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).
- [12] 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).
- [13] M. Benedetti, D. Garcia-Pintos, O. Perdomo, V. Leyton-Ortega, Y. Nam, and A. Perdomo-Ortiz, *A Generative Modeling Approach for Benchmarking and Training Shallow Quantum Circuits*, *npj Quantum Inf.* **5**, 45 (2019).
- [14] J. Alcazar, V. Leyton-Ortega, and A. Perdomo-Ortiz, *Classical Versus Quantum Models in Machine Learning: Insights from a Finance Application*, *Mach. Learn.* **1**, 035003 (2020).
- [15] M. Schuld and N. Killoran, *Quantum Machine Learning in Feature Hilbert Spaces*, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [16] X. Gao, Z.-Y. Zhang, and L.-M. Duan, *A Quantum Machine Learning Algorithm Based on Generative Models*, *Sci. Adv.* **4**, eaat9004 (2018).

- [17] 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**, 60 (2020).
- [18] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, *Expressive Power of Parametrized Quantum Circuits*, *Phys. Rev. Research* **2**, 033125 (2020).
- [19] R. Sweke, J.-P. Seifert, D. Hangleiter, and J. Eisert, *On the Quantum Versus Classical Learnability of Discrete Distributions*, *Quantum* **5**, 417 (2021).
- [20] A. Einstein, B. Podolsky, and N. Rosen, *Can Quantum-Mechanical Description of Physical Reality Be Considered Complete?*, *Phys. Rev.* **47**, 777 (1935).
- [21] J. S. Bell, *On the Einstein Podolsky Rosen Paradox*, *Phys. Phys. Fiz.* **1**, 195 (1964).
- [22] J. S. Bell, *On the Problem of Hidden Variables in Quantum Mechanics*, *Rev. Mod. Phys.* **38**, 447 (1966).
- [23] S. Kochen and E. P. Specker, *The Problem of Hidden Variables in Quantum Mechanics*, in *The Logico-Algebraic Approach to Quantum Mechanics* (Springer, New York, 1975), pp. 298–328.
- [24] J. Anders and D. E. Browne, *Computational Power of Correlations*, *Phys. Rev. Lett.* **102**, 050502 (2009).
- [25] H. Buhrman, R. Cleve, S. Massar, and R. De Wolf, *Non-locality and Communication Complexity*, *Rev. Mod. Phys.* **82**, 665 (2010).
- [26] M. Howard, J. Wallman, V. Veitch, and J. Emerson, *Contextuality Supplies the “Magic” for Quantum Computation*, *Nature (London)* **510**, 351 (2014).
- [27] J. Bermejo-Vega, N. Delfosse, D. E. Browne, C. Okay, and R. Raussendorf, *Contextuality as a Resource for Models of Quantum Computation with Qubits*, *Phys. Rev. Lett.* **119**, 120505 (2017).
- [28] M. Frembs, S. Roberts, and S. D. Bartlett, *Contextuality as a Resource for Measurement-Based Quantum Computation beyond Qubits*, *New J. Phys.* **20**, 103011 (2018).
- [29] D. Perez-Garcia, F. Verstraete, M. M. Wolf, and J. I. Cirac, *Matrix Product State Representations*, *Quantum Inf. Comput.* **7**, 401 (2007).
- [30] U. Schollwöck, *The Density-Matrix Renormalization Group in the Age of Matrix Product States*, *Ann. Phys. (Amsterdam)* **326**, 96 (2011).
- [31] E. Stoudenmire and D. J. Schwab, *Supervised Learning with Tensor Networks*, in *Advances in Neural Information Processing Systems*, edited by D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, and R. Garnett (Curran Associates, Inc., 2016), Vol. 29, pp. 4799–4807, <https://proceedings.neurips.cc/paper/2016/file/5314b9674c86e3f9d1ba25ef9bb32895-Paper.pdf>
- [32] 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).
- [33] Y. Liu, X. Zhang, M. Lewenstein, and S.-J. Ran, *Learning Architectures Based on Quantum Entanglement: A Simple Matrix Product State Algorithm for Image Recognition*, *Stat.* **1050**, 16 (2018).
- [34] I. Glasser, N. Pancotti, and J. I. Cirac, *From Probabilistic Graphical Models to Generalized Tensor Networks for Supervised Learning*, *IEEE Access* **8**, 68169 (2020).

- [35] D. Niedermayer, *An Introduction to Bayesian Networks and Their Contemporary Applications*, in *Innovations in Bayesian Networks* (Springer, New York, 2008), pp. 117–130.
- [36] G. H. Low, T. J. Yoder, and I. L. Chuang, *Quantum Inference on Bayesian Networks*, *Phys. Rev. A* **89**, 062315 (2014).
- [37] C. Manning and H. Schütze, *Foundations of Statistical Natural Language Processing* (MIT Press, Cambridge, MA, 1999).
- [38] D. M. Greenberger, M. A. Horne, A. Shimony, and A. Zeilinger, *Bell's Theorem without Inequalities*, *Am. J. Phys.* **58**, 1131 (1990).
- [39] N. Harrigan, T. Rudolph, and S. Aaronson, *Representing Probabilistic Data via Ontological Models*, [arXiv:0709.1149](https://arxiv.org/abs/0709.1149).
- [40] A. Karanjai, J. J. Wallman, and S. D. Bartlett, *Contextuality Bounds the Efficiency of Classical Simulation of Quantum Processes*, [arXiv:1802.07744](https://arxiv.org/abs/1802.07744).
- [41] R. W. Spekkens, *Contextuality for Preparations, Transformations, and Unsharp Measurements*, *Phys. Rev. A* **71**, 052108 (2005).
- [42] N. D. Mermin, *Hidden Variables and the Two Theorems of John Bell*, *Rev. Mod. Phys.* **65**, 803 (1993).
- [43] A. Peres, *Two Simple Proofs of the Kochen-Specker Theorem*, *J. Phys. A* **24**, L175 (1991).
- [44] D. Gottesman, Ph.D. thesis, California Institute of Technology, 1997.
- [45] N. S. Müller, M. Studer, and G. Ritschard, *Classification de Parcours de vie à l'Aide de l'Optimal Matching*, in *XIVe Rencontre de la Société Francophone de Classification (SFC 2007)*, pp. 157–160.
- [46] G. Ritschard, R. Bürgin, and M. Studer, *Exploratory Mining of Life Event Histories*, in *Contemporary Issues in Exploratory Data Mining in the Behavioral Sciences* (Routledge, New York, 2013), pp. 243–276.
- [47] D. Dua and C. Graff, *UCI Machine Learning Repository* (2017), <http://archive.ics.uci.edu/ml>.
- [48] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and I. Cirac, *Expressive Power of Tensor-Network Factorizations for Probabilistic Modeling*, in *Advances in Neural Information Processing Systems 32*, edited by H. Wallach *et al.* (Curran Associates, 2019), Vol. 32, pp. 1498–1510, <https://proceedings.neurips.cc/paper/2019/file/b86e8d03fe992d1b0e19656875ee557c-Paper.pdf>.
- [49] H. Jaeger, *Observable Operator Models for Discrete Stochastic Time Series*, *Neural Comput.* **12**, 1371 (2000).
- [50] M.-J. Zhao and H. Jaeger, *Norm-Observable Operator Models*, *Neural Comput.* **22**, 1927 (2010).
- [51] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, *Barren Plateaus in Quantum Neural Network Training Landscapes*, *Nat. Commun.* **9**, 4812 (2018).
- [52] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, *Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks*, *Nat. Commun.* **12**, 1791 (2021).
- [53] C. O. Marrero, M. Kieferová, and N. Wiebe, *Entanglement Induced Barren Plateaus*, [arXiv:2010.15968](https://arxiv.org/abs/2010.15968).
- [54] T. L. Patti, K. Najafi, X. Gao, and S. F. Yelin, *Entanglement Devised Barren Plateau Mitigation*, *Phys. Rev. Research* **3**, 033090 (2021).
- [55] Z. Holmes, A. Arrasmith, B. Yan, P. J. Coles, A. Albrecht, and A. T. Sornborger, *Barren Plateaus Preclude Learning Scramblers*, *Phys. Rev. Lett.* **126**, 190501 (2021).
- [56] E. R. Anschuetz, *Critical Points in Hamiltonian Agnostic Variational Quantum Algorithms*, [arXiv:2109.06957](https://arxiv.org/abs/2109.06957).
- [57] K. Poland, K. Beer, and T. J. Osborne, *No Free Lunch for Quantum Machine Learning*, [arXiv:2003.14103](https://arxiv.org/abs/2003.14103).
- [58] M. S. Rudolph *et al.*, *Generation of High-Resolution Handwritten Digits with an Ion-Trap Quantum Computer*, [arXiv:2012.03924](https://arxiv.org/abs/2012.03924).
- [59] D. Poulin, A. Qarry, R. Somma, and F. Verstraete, *Quantum Simulation of Time-Dependent Hamiltonians and the Convenient Illusion of Hilbert Space*, *Phys. Rev. Lett.* **106**, 170501 (2011).
- [60] R. Orús, *A Practical Introduction to Tensor Networks: Matrix Product States and Projected Entangled Pair States*, *Ann. Phys. (Amsterdam)* **349**, 117 (2014).
- [61] J. C. Bridgeman and C. T. Chubb, *Hand-Waving and Interpretive Dance: An Introductory Course on Tensor Networks*, *J. Phys. A* **50**, 223001 (2017).
- [62] F. Verstraete, V. Murg, and J. I. Cirac, *Matrix Product States, Projected Entangled Pair States, and Variational Renormalization Group Methods for Quantum Spin Systems*, *Adv. Phys.* **57**, 143 (2008).
- [63] G. Vidal, *Entanglement Renormalization*, *Phys. Rev. Lett.* **99**, 220405 (2007).
- [64] F. Verstraete and J. I. Cirac, *Renormalization Algorithms for Quantum-Many Body Systems in Two and Higher Dimensions*, [arXiv:cond-mat/0407066](https://arxiv.org/abs/cond-mat/0407066).
- [65] J. H. Martin and D. Jurafsky, *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition* (Pearson/Prentice Hall, Englewood Cliffs, NJ, 2009).
- [66] G. E. Hinton, *Deep Belief Networks*, *Scholarpedia* **4**, 5947 (2009).
- [67] V. Bergholm, J. J. Vartiainen, M. Möttönen, and M. M. Salomaa, *Quantum Circuits with Uniformly Controlled One-Qubit Gates*, *Phys. Rev. A* **71**, 052330 (2005).
- [68] S. Kullback and R. A. Leibler, *On Information and Sufficiency*, *Ann. Math. Stat.* **22**, 79 (1951).
- [69] N. D. Mermin, *Extreme Quantum Entanglement in a Superposition of Macroscopically Distinct States*, *Phys. Rev. Lett.* **65**, 1838 (1990).
- [70] R. Raussendorf and H. J. Briegel, *One-Way Quantum Computer*, *Phys. Rev. Lett.* **86**, 5188 (2001).
- [71] J. Barrett, C. M. Caves, B. Eastin, M. B. Elliott, and S. Pironio, *Modeling Pauli Measurements on Graph States with Nearest-Neighbor Classical Communication*, *Phys. Rev. A* **75**, 012103 (2007).
- [72] S. Bravyi, D. Gosset, and R. Koenig, *Quantum Advantage with Shallow Circuits*, *Science* **362**, 308 (2018).
- [73] A. Peres, *Incompatible Results of Quantum Measurements*, *Phys. Lett. A* **151**, 107 (1990).
- [74] N. D. Mermin, *Simple Unified Form for the Major No-Hidden-Variables Theorems*, *Phys. Rev. Lett.* **65**, 3373 (1990).
- [75] P. Aravind, *A Simple Demonstration of Bell's Theorem Involving Two Observers and No Probabilities or Inequalities*, [arXiv:quant-ph/0206070](https://arxiv.org/abs/quant-ph/0206070).
- [76] S. Aaronson and D. Gottesman, *Improved Simulation of Stabilizer Circuits*, *Phys. Rev. A* **70**, 052328 (2004).
- [77] K. B. Petersen and M. S. Pedersen, *The Matrix Cookbook* (2012), http://www2.imm.dtu.dk/pubdb/views/publication_details.php?id=3274, Version 20121115.

- [78] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, *A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains*, *Ann. Math. Stat.* **41**, 164 (1970).
- [79] D. M. Chickering, *Learning Bayesian Networks Is NP-Complete*, in *Learning from Data* (Springer, New York, 1996), pp. 121–130.
- [80] P. Dagum and M. Luby, *Approximating Probabilistic Inference in Bayesian Belief Networks Is NP-Hard*, *Artif. Intell.* **60**, 141 (1993).
- [81] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, Cambridge, MA, 2009).
- [82] T. E. Abrudan, J. Eriksson, and V. Koivunen, *Steepest Descent Algorithms for Optimization under Unitary Matrix Constraint*, *IEEE Trans. Signal Proc.* **56**, 1134 (2008).
- [83] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Learning Representations by Back-Propagating Errors*, *Nature (London)* **323**, 533 (1986).
- [84] B. Coecke, G. de Felice, K. Meichanetzidis, and A. Toumi, *Foundations for Near-Term Quantum Natural Language Processing*, *arXiv:2012.03755*.
- [85] C. Guo, Z. Jie, W. Lu, and D. Poletti, *Matrix Product Operators for Sequence to Sequence Learning*, *Phys. Rev. E* **98**, 042114 (2018).
- [86] E. M. Stoudenmire, *Learning Relevant Features of Data with Multi-Scale Tensor Networks*, *Quantum Sci. Technol.* **3**, 034003 (2018).
- [87] K. Temme and F. Verstraete, *Stochastic Matrix Product States*, *Phys. Rev. Lett.* **104**, 210502 (2010).
- [88] L. Isenhower, M. Saffman, and K. Mølmer, *Multibit C_k NOT Quantum Gates via Rydberg Blockade*, *Quantum Inf. Process.* **10**, 755 (2011).
- [89] P. Høyer and R. Špalek, *Quantum Fan-out Is Powerful*, *Theory Comput.* **1**, 81 (2005).
- [90] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, Cambridge, England, 2009).
- [91] A. W. Harrow and A. Montanaro, *Quantum Computational Supremacy*, *Nature (London)* **549**, 203 (2017).
- [92] X. Gao, S.-T. Wang, and L.-M. Duan, *Quantum Supremacy for Simulating a Translation-Invariant Ising Spin Model*, *Phys. Rev. Lett.* **118**, 040502 (2017).
- [93] L. Stockmeyer, *On Approximation Algorithms for #P*, *SIAM J. Comput.* **14**, 849 (1985).
- [94] R. M. Karp and R. Lipton, *Turing Machines that Take Advice*, *Enseign. Math.* **28**, 191 (1982).
- [95] S. Aaronson, A. Cojocaru, A. Gheorghiu, and E. Kashefi, *On the Implausibility of Classical Client Blind Quantum Computing*, *arXiv:1704.08482*.
- [96] H. Larochelle and Y. Bengio, *Classification Using Discriminative Restricted Boltzmann Machines*, in *Proceedings of the 25th International Conference on Machine Learning* (Association for Computing Machinery, New York, 2008), pp. 536–543, [10.1145/1390156.1390224](https://doi.org/10.1145/1390156.1390224).
- [97] M. Nest, *Classical Simulation of Quantum Computation, the Gottesman-Knill Theorem, and Slightly Beyond*, *Quantum Inf. Comput.* **10**, 0258 (2010).
- [98] M. Krishna, Ph.D. thesis, University of Kerala, 2009 (<http://hdl.handle.net/10603/94513>).