

**Differentiable learning of quantum circuit Born machines**Jin-Guo Liu<sup>1</sup> and Lei Wang<sup>1,2,\*</sup><sup>1</sup>*Institute of Physics, Chinese Academy of Sciences, Beijing 100190, China*<sup>2</sup>*CAS Center for Excellence in Topological Quantum Computation, University of Chinese Academy of Sciences, Beijing 100190, China* (Received 16 April 2018; revised manuscript received 24 September 2018; published 19 December 2018)

Quantum circuit Born machines are generative models which represent the probability distribution of classical dataset as quantum pure states. Computational complexity considerations of the quantum sampling problem suggest that the quantum circuits exhibit stronger expressibility compared to classical neural networks. One can efficiently draw samples from the quantum circuits via projective measurements on qubits. However, similar to the leading implicit generative models in deep learning, such as the generative adversarial networks, the quantum circuits cannot provide the likelihood of the generated samples, which poses a challenge to the training. We devise an efficient gradient-based learning algorithm for the quantum circuit Born machine by minimizing the kernelized maximum mean discrepancy loss. We simulated generative modeling of the BARS-AND-STRIPES dataset and Gaussian mixture distributions using deep quantum circuits. Our experiments show the importance of circuit depth and the gradient-based optimization algorithm. The proposed learning algorithm is runnable on near-term quantum device and can exhibit quantum advantages for probabilistic generative modeling.

DOI: [10.1103/PhysRevA.98.062324](https://doi.org/10.1103/PhysRevA.98.062324)**I. INTRODUCTION**

Unsupervised generative modeling is at the forefront of deep learning research [1]. Unlike the extremely successful discriminative tasks, such as supervised classification and regression, the goal of generative modeling is to model the probability distribution of observed data and generate new samples accordingly. Generative modeling finds wide applications in computer vision [2], speech synthesis [3], as well as chemical design [4]. And it is believed to be a crucial component towards artificial general intelligence. However, generative modeling is more challenging than discriminative tasks since it requires one to efficiently represent, learn, and sample from high-dimensional probability distributions [1].

In parallel to the rapid development of deep learning, there is heated ongoing research to fabricate intermediate-scale quantum circuits [5–7]. Quantum hardware may show greater potential than their classical counterparts in generative tasks. Multiple schemes have been proposed to boost the performance of classical generative models using quantum devices. For example, quantum Boltzmann machines [8–11] generalize the energy function of classical Boltzmann machines [12] to quantum Hamiltonian for possible stronger representational power and faster training. A concern which may prevent the quantum Boltzmann machine from surpassing its classical counterpart is the limited connectivity on the actual quantum hardware [13]. Reference [14] introduces a quantum generalization of the probabilistic graphical model [15], which can be exponentially more powerful than its classical counterpart and has exponential speedup in training and inference, at least, for some instances under reasonable assumptions in computational complexity theory.

Another class of arguably simpler quantum generative models, named Born machines [16–18], directly exploit the inherent probabilistic interpretation of quantum wave functions [19]. Born machines represent probability distribution using a quantum pure state instead of the thermal distribution, such as the Boltzmann machines [12]. Therefore, Born machines can directly generate samples via projective measurement on the qubits, in contrast to the slow mixing Gibbs sampling approach. Moreover, computational complexity considerations on quantum sampling problems suggest that a quantum circuit can produce probability distribution that is P-hard [20–23], which is infeasible to simulate efficiently using classical algorithms. The same reasoning underlines the current efforts towards “quantum supremacy” experiments by sampling outcomes of random quantum circuits [24]. Reference [16] performed a classical simulation of the Born machine using the matrix product state representation of the quantum state. It will be even more promising to realize the Born machines using quantum circuits since one can, at least, efficiently prepare some of the tensor networks on a quantum computer [25]. Recently, Ref. [18] demonstrated experimental realization of the Born machine on a four-qubit shallow quantum circuit trained by gradient-free optimization of measurement histograms.

To further scale up the quantum circuit Born machine (QCBM) to larger numbers of qubits and circuit depth, one needs to devise an appropriate objective function for the generative tasks without explicit reference to the model probability. Unlike the tensor network simulation [16], QCBM belongs to *implicit* generative models since one does not have access to the wave function of an actual quantum circuit. Thus, the QCBM can be used as a simulator to generate samples without access to their likelihoods, which is similar to the notable generative adversarial networks (GANs) [26,27]. Compared to generative models with explicit likelihoods, such

\*wanglei@iphy.ac.cn

as the Boltzmann machines [28], normalizing flows [29–33], and variational autoencoders [34,35], the implicit generative models can be more expressive due to less restrictions in their network structures. On the other hand, having no direct access to the output probability also poses challenges to the scalable training of quantum circuits.

Moreover, one also needs better learning algorithm than the gradient-free optimization scheme [18] especially given the noisy realization of current quantum circuits. Similarly, scalability of the optimization scheme is also a crucial concern in deep learning in which deep neural networks can even reach billions of parameters [36]. In the history of machine learning, gradient-free algorithms were employed to optimize small-scale neural networks [37]. However, they failed to scale up to a larger number of parameters. It is the back-propagation algorithm [38] which can efficiently compute the gradient of the neural network output with respect to the network parameters enables scalable training of deep neural nets. It is thus highly demanded to have scalable quantum algorithms [39,40] for estimating gradients on actual quantum circuits.

Recently, gradient-based learning of quantum circuits has been devised for quantum control [41] and discriminative tasks [42,43]. Although they are still less efficient compared to the back-propagation algorithm for neural networks, these unbiased gradient algorithms can already greatly accelerate the quantum circuit learning. Unfortunately, direct application of these gradient algorithms [41–43] to QCBM training is still nontrivial since the output of the generative model is genuinely bit strings which follow high-dimensional probability distributions. The previous method to get the gradient requires the target function being an expectation value of some observable, this restriction can make almost every meaningful loss function in generative modeling fail to differentiate since even simplest mean-square loss (with respect to probabilities) contains exponentially many observables which makes it impossible to implement. In fact, it is even an ongoing research topic in deep learning to perform differentiable learning of an implicit generative model with discrete outputs [27,44].

In this paper, we develop an efficient gradient-based learning algorithm to train the QCBM. In what follows, we first present a practical quantum-classical hybrid algorithm to train the quantum circuit as a generative model in Sec. II, thus realizes a Born machine. Then, we apply the algorithm on  $3 \times 3$  BARS-AND-STRIPES and double Gaussian peak datasets in Sec. III. We show that the training is robust to moderate sampling noise and is scalable in circuit depth. Increasing the circuit depth significantly improves the representational power for generative tasks. Finally, we conclude and discuss caveats and future research directions about the QCBM in Sec. IV.

## II. MODEL AND LEARNING ALGORITHM

Given a dataset  $\mathcal{D} = \{x\}$  containing independent and identically distributed (i.i.d.) samples from a target distribution  $\pi(x)$ , we set up a QCBM to generate samples close to the unknown target distribution. As shown in Fig. 1, the QCBM takes the product state  $|0\rangle$  as an input and evolves it to a final-state  $|\psi_\theta\rangle$  by a sequence of unitary gates. Then we can measure this output state on a computation basis to obtain a

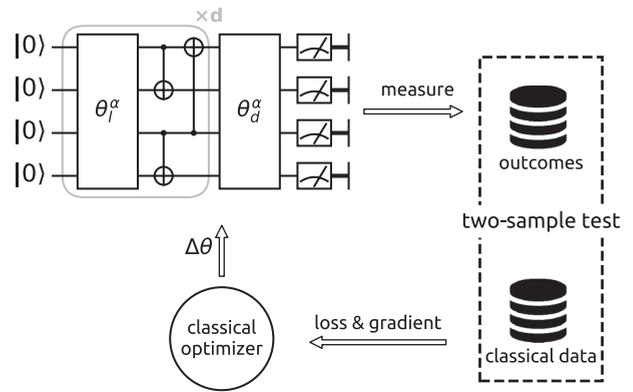


FIG. 1. Illustration of the differentiable QCBM training scheme. The top left is the quantum circuit which produces bit string samples. The dashed box on the right denotes the two-sample test on the generated samples and training samples with the loss function [Eq. (1)] and corresponding gradients [Eq. (2)] as outputs.  $\Delta\theta$  is the amount of update applied to the circuit parameters, which are computed by a classical optimizer. The outcome of the training is to produce a quantum circuit which generates samples according to the learned probability distribution on the computational basis.

sample of bits  $x \sim p_\theta(x) = |\langle x|\psi_\theta\rangle|^2$ . The goal of the training is to let the model probability distribution  $p_\theta$  approach  $\pi$ .

We employ a classical-quantum hybrid feedback loop as the training strategy. The setup is similar to the quantum approximate optimization algorithm [45–47] and the variational quantum eigensolver [48–50]. By constructing the circuits and performing measurements repeatedly, we collect a batch of samples from the QCBM. Then we introduce two-sample test as a measure of distance between the generated samples and the training set, which is used as our differentiable loss. Using a classical optimizer which takes the gradient information of the loss function, we can push the generated sample distribution towards the target distribution.

### A. Quantum circuit architecture design

The overall circuit layout is similar to the IBM variational quantum eigensolver [50] where one interweaves single qubit rotation layers and entangler layers shown in Fig. 1. The rotation layers are parametrized by rotation angles  $\theta = \{\theta_l^\alpha\}$  where the layer index  $l$  runs from 0 to  $d$  with  $d$  as the maximum depth of the circuit.  $\alpha$  is a combination of qubit index  $j$  and an arbitrary rotation gate index where the arbitrary rotation gate has the form  $U(\theta_l^j) = R_z(\theta_l^{j,1})R_x(\theta_l^{j,2})R_z(\theta_l^{j,3})$  with  $R_m(\theta) \equiv \exp(-\frac{i\theta\sigma_m}{2})$ . The total number of parameters in this QCBM is  $(3d + 1)n$  with  $n$  as the number of qubits [51].

We employ CNOT gates with no learnable parameters for the entangle layers to induce correlations between qubits. In light of experimental constraints on the connectivity of the circuits, we make the connection of the entangle layers to be sparse by requiring its topology as a tree (i.e., the simplest connected graph). From the classical probabilistic graphical model’s perspective [15], the tree graph that captures information content of the dataset most efficiently is the Chow-Liu tree [52]. Since controlled unitary gates have a close relation with classical probability graphical models [53],

we employ the same Chow-Liu tree as the topology of CNOT gates. To construct the Chow-Liu tree, we first compute mutual information between all pairs of the bits for samples in the training set as weights and then construct the maximum spanning tree using, for example, Kruskal's algorithm. The assignment of the control bit and the target bit on a bond is random since the Chow-Liu algorithm treated directed and undirected graphs the same. In the case where this connection structure is not directly supported by the hardware, a combination of SWAP gates and CNOT gates can be used to efficiently simulate the required structure [54].

The performance of entangle layers constructed in this procedure is better than most random connections with the same number of gates. This data-driven quantum circuit architecture design scheme that puts more entanglers between highly correlated bits, captures the correlation between qubits better, and may alleviate issues of vanishing gradients for large-scale applications [25].

### B. Loss function and gradient-based optimization

Viewing the QCBM as an implicit generative model [55], we train it by employing the kernel two-sample test [56]. The idea is to compare the distance in the kernel feature space on the samples drawn from the target and the model distributions. We refer the following loss function as the squared maximum mean discrepancy (MMD) [57,58]:

$$\begin{aligned} \mathcal{L} &= \left\| \sum_x p_\theta(x) \phi(x) - \sum_x \pi(x) \phi(x) \right\|^2 \\ &= \mathbb{E}_{x \sim p_\theta, y \sim p_\theta} [K(x, y)] - 2 \mathbb{E}_{x \sim p_\theta, y \sim \pi} [K(x, y)] \\ &\quad + \mathbb{E}_{x \sim \pi, y \sim \pi} [K(x, y)]. \end{aligned} \quad (1)$$

The summation in the first line runs over the whole Hilbert space. The expectation values in the second line are for the corresponding probability distributions. The function  $\phi$  maps  $x$  to a high-dimensional reproducing kernel Hilbert space [59]. However, as common in the kernel methods, by defining a kernel function  $K(x, y) = \phi(x)^T \phi(y)$ , one can avoid working in the high-dimensional feature space. We employ a Gaussian kernel  $K(x, y) = \exp(-\frac{1}{2\sigma^2} |x - y|^2)$  to reveal differences between the two distributions under various scales. Here,  $\sigma$  is the bandwidth parameter which controls the width of the Gaussian kernel. The sample  $x$  can either be a bit string (vector) or be an integer (scalar) depending on the representation. When  $x$  is a bit string,  $|x|$  stands for the  $\ell^2$  norm in the vector space. The MMD loss with Gaussian kernels asymptotically approaches zero if and only if the output distribution matches the target distribution exactly [57,58]. The same loss function was used to train the generative moment matching networks (GMMNs) [60–62].

Previous methods to get gradients with respect to circuit parameters in Ref. [43] requires the target function being an observable. However, a statistic functional, such as our MMD loss cannot be expressed as a combination of finite (i.e., not exponentially many) observables. To learn the QCBM as a generative model, we formulated the unbiased gradient of

MMD loss Eq. (1) in Appendix A as

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta_l^\alpha} &= \mathbb{E}_{x \sim p_{\theta^+}, y \sim p_\theta} [K(x, y)] \\ &\quad - \mathbb{E}_{x \sim p_{\theta^-}, y \sim p_\theta} [K(x, y)] \\ &\quad - \mathbb{E}_{x \sim p_{\theta^+}, y \sim \pi} [K(x, y)] \\ &\quad + \mathbb{E}_{x \sim p_{\theta^-}, y \sim \pi} [K(x, y)]. \end{aligned} \quad (2)$$

Here,  $p_{\theta^+}(x)$  and  $p_{\theta^-}(x)$  are output probabilities of the QCBM under circuit parameters  $\theta^\pm = \theta \pm \frac{\pi}{2} \mathbf{e}_l^\alpha$ , where  $\mathbf{e}_l^\alpha$  is the  $(l, \alpha)$ th unit vector in parameter space (i.e.,  $\theta_l^\alpha \leftarrow \theta_l^\alpha \pm \frac{\pi}{2}$ , with other angles unchanged). In contrast to the finite difference methods, such as simultaneous perturbation stochastic approximation (SPSA) [63], Eq. (2) is an unbiased estimator of the exact gradient.

In order to estimate the gradient [Eq. (2)] on an actual quantum circuit, one can repeatedly send rotation and entangle pulses to the device according to the circuit parameters  $\theta^{(\pm)}$  and then perform projective measurements on the computational basis to collect binary samples  $x \sim p_{\theta^{(\pm)}}$ . Whereas for  $x \sim \pi$ , one can simply take a batch of data from the training dataset [64]. The sampling noise in the estimated gradient is controlled by the number of measurements  $N$ , denoted as the batch size. After one has obtained a sufficiently accurate gradient, one can use a classical optimizer to update the circuit parameters similar to the stochastic gradient descent training of deep neural nets.

Parameter learning of quantum circuits is adaptive in the sense that the implementation of quantum gates can even be nonideal. One can obtain gradients with high accuracy as long as the parametrized single-qubit rotation gates are precise, which is relatively easier to achieve experimentally. The optimization scheme is independent of the detailed form of nonparametrized entangle gates. Thus, the CNOT gate in the setup can be replaced by any gate which can generate desired quantum entanglements.

It is instructive to compare the training of the QCBM to that of classical implicit generative models, such as the GAN [26,27] and the GMMN [60–62]. Classically, one does not have access to the likelihood either. The gradient is thus obtained via the chain rule  $\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}}{\partial x} \frac{\partial x}{\partial \theta}$ , which then does not apply for discrete data. On the other hand, the unbiased gradient estimator of the QCBM takes advantage of the known structure of the unitary evolution and the MMD loss (see Appendix A), despite that the probability of the outcome is unknown. In this sense, quantum circuits exhibit a clear quantum advantage over classical neural nets since they fill the gap of differentiable learning of implicit generative models of discrete data.

### III. NUMERICAL EXPERIMENTS

We carry out numerical experiments by simulating the learning of the QCBM on a classical computer. These experiments reveal advantages of gradient based optimization over gradient-free optimization and demonstrate stronger expressibility of deep circuits over shallow ones. The code can be

found in Ref. [65], and the corresponding julia version can be found in Ref. [66].

**A. Bars-and-stripes dataset**

We first train a QCBM on the BARS-AND-STRIPES dataset [16,18,67], which is a prototypical image dataset consisting of vertical bars and horizontal stripes. On a grid of  $3 \times 3$ , the dataset contains 14 valid configurations. We model the pixels with a quantum circuit of nine qubits. The Chow-Liu tree for this dataset is shown in Fig. 2(a). Bonds are either rowwise or columnwise since correlations of pixels sharing the same row or column index are dominant in this dataset. The bandwidth used in Gaussian kernels of the MMD loss is  $\sigma = 2$ .

For circuit depth  $d = 10$ , our gradient-based training is able to reduce the MMD loss efficiently. The loss function for different iteration steps is shown in Fig. 3(a). We first perform L-BFGS-B [68] optimization (black dashed line) using the exact gradient computed via the wave function ( $N = \infty$ ) to test the expressibility of the quantum circuit. A loss of  $3.3 \times 10^{-6}$  can be achieved, showing that the circuit is quite expressive in terms of the two-sample test.

In practice, one has to perform projective measurements on the qubits to collect statistics of the gradient since the wave function is inaccessible. This situation is similar to the minibatch estimate of the gradient in deep learning [1]. As is well known in the deep learning applications [1], the L-BFGS-B algorithm is not noise tolerant. Thus, it is unsuitable for quantum circuit learning in a realistic situation. One needs to employ an alternate optimizer which is robust to the sampling noise to train the quantum circuit with the noisy gradient estimator.

We employ the stochastic gradient optimizer Adam [69] with the learning rate 0.1. The sampling noise in the gradients can be controlled by tuning the batch size  $N = 2000, 20\,000, \infty$  of the measurements. The solid lines in Fig. 3(a) show that, as the sample size increases, the final MMD loss reduces systematically. The scatters in the inset confirmed that the model probability of the learned quantum

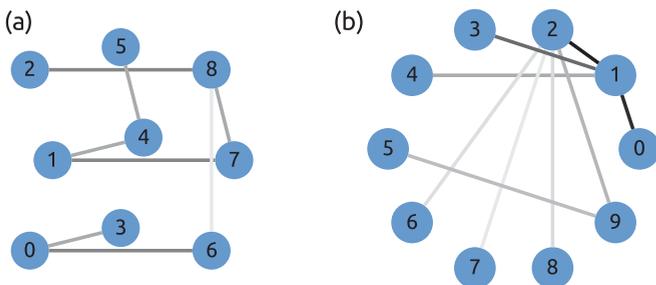


FIG. 2. (a) Connectivity of the CNOT gates for the  $3 \times 3$  BARS-AND-STRIPES dataset generated via the Chow-Liu tree algorithm. The qubits are arranged on a  $3 \times 3$  grid with some of them shifted a bit in order to visualize the edges clearly. (b) The Chow-Liu tree for the double-Gaussian peak model; the numbers represent the position of a bit in the digit, 0 for the big end and 9 for the little end. In this plot, the darkness of the edges indicates the amount of mutual information between two sites.

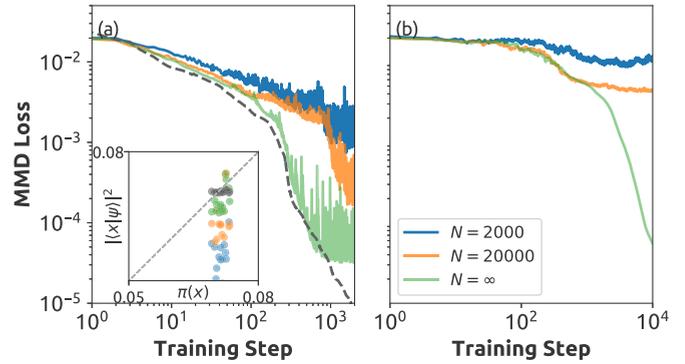


FIG. 3. The MMD loss [Eq. (1)] as a function of training steps under different sampling errors that governed by batch size  $N$ . (a) Gradient-based training, solid colored lines are for Adam, and the dashed black line is for Limited-memory Broyden-Fletcher-Goldfarb-Shanno with Box constraints (L-BFGS-B) with  $N = \infty$ . The inset is a comparison of probability distribution between the training set and the QCBM output; the points on the dashed line mean exact matches. (b) The gradient-free CMA-ES training counterpart where each point on the graph represents a mean loss of its population.

circuit and the target probability aligns better with the lower MMD loss.

To visualize the quality of the samples, we generated a few samples from the QCBM trained under different measurement batch size  $N$  in Fig. 4. Here, we define a valid rate  $\chi \equiv p(x \text{ is a bar or a stripe})$  as a measure of generation quality. The valid rate increases as the batch size increases. However, even with a moderate number of measurement  $N = 2000$ , one can achieve a valid rate of  $\chi = 74.4\%$ . Here, we should mention that the best valid rate of a  $d = 10$  layer circuit is



FIG. 4.  $3 \times 3$  BARS-AND-STRIPES samples generated from the QCBMs. The circuit parameters used here are from the final stages of Adam training with different batch sizes  $N$  in Fig. 3(a).  $\chi$  is the rate of generating valid samples in the training dataset. For illustrative purposes, we only show 12 samples for each situation with batch size  $N$ .

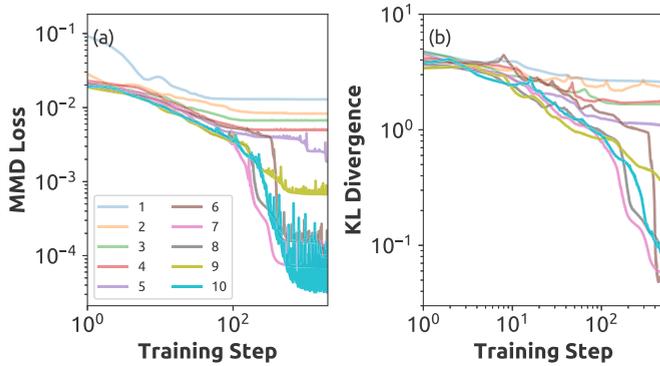


FIG. 5. Losses as a function of training step for circuit depth  $d = 1, \dots, 10$  (the darker the color of a line, the deeper the circuit). (a) The MMD loss Eq. (1) and (b) the corresponding Kullback-Leibler (KL) divergence. Here, we use the Adam optimizer with an exact gradient.

achieved by the L-BFGS-B optimizer with  $N = \infty$ , which is  $\chi = 99.0\%$ .

To highlight the importance of using a gradient-based optimizer, we compare our approach to the covariance matrix adaptation evolution strategy (CMA-ES) [70,71], a state-of-the-art gradient-free stochastic optimizer. The input of the CMA-ES is the scalar loss function measured on the circuit instead of the vector gradient information. The CMA-ES optimizer is able to optimize nonsmooth nonconvex loss functions efficiently [71] and thus, in general, performs better than other gradient-free methods, such as the SPSA [63] in training noisy quantum circuits. We have confirmed this in our simulation.

In the absence of sampling noise  $N = \infty$  in Fig. 3(b), we do observe that the CMA-ES optimizer is able to achieve a similar performance as the Adam optimizer after  $10^4$  steps of optimization with a population size of 50. The total number of generated samples is  $10^4 \times 50 \times N$ , which is comparable to the Adam training in Fig. 3(a) [72].

However, the performance of the CMA-ES deteriorates significantly once taking sampling noise into consideration as shown for  $N = 2000$  and  $N = 20000$  in Fig. 3(b). A possible explanation is that, in each step of the CMA-ES, its evolution strategy chooses the direction to go by inspecting the center of top 20% instances. This process can be understood as an effective finite difference gradient estimation based on the losses of its population. However, extracting gradient information from noisy losses is difficult, even when one has plenty of them.

Another advantage of using gradient-based learning is the efficiency comparing with gradient-free methods, which gets particularly significant when circuits get deeper and the number of parameters increases. In the following, we address the necessity for using deep circuits. Figure 5(a) shows the MMD loss as a function of Adam training steps for different circuit depths. One obtains lower loss for deeper quantum circuit after 2000 optimization steps. Figure 5(b) shows the KL divergence [73] calculated using the circuit parameters in (a) at different training steps. Note that this quantity is inaccessible for large-scale problems since one has no access to the target nor the output probability. We compute the KL divergence for the toy model to demonstrate that the MMD loss is a good

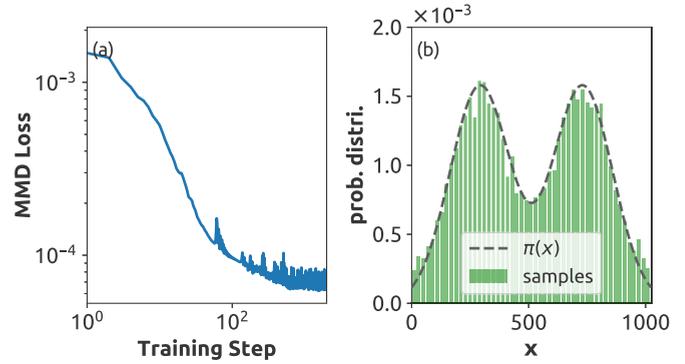


FIG. 6. (a) The MMD loss as a function of the Adam training step. (b) Histogram for samples generated by a trained QCBM with a bin width of 20 (green bars) in comparison with the exact probability density function (black dashed line).

surrogate for practical training. The result, indeed, shows a consistency between MMD loss and KL divergence. And it also supports the observation that deep circuits have stronger representational power. Similar to deep neural networks, deep circuits can achieve better performance also due to that one is less prone to be trapped in a poor local minima with larger amounts of parameters [74].

Another advantage of the QCBM over traditional deep neural networks is that its training does not suffer from gradient vanishing or the exploding problem as the circuit goes deeper. Gradient vanishing or exploding is a common problem for a traditional deep neural network [75] which originates from multiplications of a long chain of matrices in the back-propagation algorithm. Training of the deep quantum circuits naturally circumvented this problem due to the unitary property of the time evolution. A similar idea was exploited in constructing classical recurrent neural networks with unitary building blocks [76]. More numerical simulation and analytical explanations can be found in Appendix B.

## B. Mixture of Gaussians

Next, we train a QCBM to model a mixture of Gaussian distributions,

$$\pi(x) \propto e^{-(1/2)[(x-\mu_1)/v]^2} + e^{-(1/2)[(x-\mu_2)/v]^2}. \quad (3)$$

Here,  $x = 1, \dots, x_{\max}$  is an integer encoded by the qubits with  $x_{\max} = 2^n$  and  $n$  is the number of qubits. It is different from the BARS-AND-STRIPES dataset in which case a sample  $x$  is represented as a bit string. We choose  $v = \frac{1}{8}x_{\max}$ , the centers  $\mu_1 = \frac{2}{7}x_{\max}$ , and  $\mu_2 = \frac{5}{7}x_{\max}$ . The distribution is shown as the dashed line in Fig. 6(b).

In the following discussion, we use  $n = 10$  qubits and set circuit depth  $d = 10$ . Unlike the case of BARS-AND-STRIPES, the Gaussian mixture distribution is smooth and nonzero for all basis states. Here, we generate  $10^5$  i.i.d. samples from the target distribution as the training set. Its Chow-Liu tree is shown in Fig. 2(b). In this graph, we see the main contributions of mutual information are from bits near the big end (most significant bits labeled by small indices). This is because the bit near the little end only determines the local translation of the probability on data axis. But for

a smooth probability distribution, the value of the little end is nearly independent from values of the rest bits. For example, the value of the big-end 0th bit being 0/1 corresponds to the global left or right peak in Fig. 6(b). Whereas the probability for the little end being 0/1 corresponds to  $x$  being even or odd.

Figure 6(a) shows the MMD loss as a function of the Adam optimization steps with a sample size of  $N = \infty$ . Here, the bandwidth of the Gaussian kernel in the MMD loss is  $\sigma = 2$ .

After 2000 training steps, the MMD loss decreased from  $1.8 \times 10^{-3}$  to  $7 \times 10^{-5}$ . To see whether this low MMD loss represents a good generative model, we then generate 20 000 samples from the QCBM and plot its binned histogram in Fig. 6(b). We see an excellent match between the histogram (green bars) and the exact probability distribution (black dashed curve). Thus, we conclude that the MMD loss with the Adam optimizer can also learn a smooth probability distribution over the qubit index of a QCBM. Here, we acknowledge that the unbinned histogram appears more spiky, partly due to the MMD loss not capturing the local variation of the probability distribution. Better circuit architecture design for representing continuous distribution may help alleviate this problem.

#### IV. DISCUSSIONS

We presented a practical gradient-based learning scheme to train quantum circuit Born machine as a generative model of discrete data. With a hardware efficient quantum circuit, the key component of the learning algorithm is to measure the gradient of the MMD two-sample test loss function Eq. (2) on a quantum computer unbiasedly and efficiently. Quantum circuits turn out to be very suited for generative modeling with discrete latent space. As training generative models with discrete output in a classical setup cannot be easily performed (usually, we need inefficient one-shot encoding), it provides a strong motivation to embrace the quantum world.

Our results exemplify the extensibility and noise tolerance of our learning scheme, which underlines the experimental achievability for the quantum advantage. With the ability to scale up, we simulated the training of intermediate-size deep circuits. Results show strong evidence that building deeper circuits will increase the representation power consistently without suffering from the notorious gradient vanishing or exploding problem in classical deep learning. Our paper also inspires an approach for quantum state preparation. By changing measuring bases randomly, we can learn how to generate target a quantum wave function by fitting measuring outcomes in every basis.

A recent work [77] pointed out that the quantum circuit also faces a gradient vanishing problem as the number of qubits increases. They found that the variance of gradient amplitudes decreases exponentially as the number of qubits increases in a random quantum circuit. However, it is reasonable to believe that with a better circuit structure design and a parametrization strategy, such as in Sec. II A and Ref. [25], or using shared weights, such as performed in the convolutional neural networks [1], the gradient vanishing problem can be alleviated. How gradient vanishing really affects gradient-based training in a large-scale QCBM needs further systematic investigation.

Our simulation of the quantum circuits Born machine is limited to a small number of qubits, thus the training set contains all patterns, and we are pushing the circuits towards the memorization limit. In future applications, one will focus on the generalization ability of the quantum circuits. In those cases, the structure and depth of a quantum circuit provide the means of regularization since they can be designed to express inductive bias in the natural distributions. In terms of the learning, the randomness in the stochastic gradient is also in favor of the generalization [74].

Besides the two-sample test loss employed in this paper, one can explore alternative training schemes, e.g., adversarial training [26]. Alternatively, learning of the kernel function used in Eq. (1) may also improve the generation quality [62]. Finally, differentiable learning of the QCBM may be used for solving combinatorial optimization problems and structure learning tasks where the outputs are encoded in discrete bit strings.

#### ACKNOWLEDGMENTS

Simulations of quantum circuits were mainly performed using the PROJECTQ library [78,79]. We thank A. Perdomo-Ortiz, M. Stoudenmire, D. Steiger, X. Gao, and P. Zhang for helpful discussions. The authors were supported by the National Natural Science Foundation of China under Grant No. 11774398 and research program of the Chinese Academy of Sciences under Grant No. XDPB0803.

#### APPENDIX A: UNBIASED GRADIENT ESTIMATOR OF THE PROBABILITY OF A QUANTUM CIRCUIT

##### 1. MMD loss

In the following, we derive Eq. (2) in the main text, starting from the partial derivative of Eq. (1),

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta_l^\alpha} &= \sum_{x,y} K(x,y) \left( p_\theta(y) \frac{\partial p_\theta(x)}{\partial \theta_l^\alpha} + p_\theta(x) \frac{\partial p_\theta(y)}{\partial \theta_l^\alpha} \right) \\ &\quad - 2 \sum_{x,y} K(x,y) \frac{\partial p_\theta(x)}{\partial \theta_l^\alpha} \pi(y). \end{aligned} \quad (\text{A1})$$

The partial derivative on the right-hand side of this equation can be unbiasedly computed using the approach of Refs. [41,43]. For a circuit containing a unitary gate parametrized as  $U(\eta) = e^{-i(\eta/2)\Sigma}$  with  $\Sigma^2 = 1$  (e.g.,  $\Sigma$  can be Pauli operators, CNOT, or SWAP), the gradient of the expected value of an observable  $B$  with respect to the parameter  $\eta$  reads

$$\frac{\partial \langle B \rangle_\eta}{\partial \eta} = \frac{1}{2} (\langle B \rangle_{\eta^+} - \langle B \rangle_{\eta^-}), \quad (\text{A2})$$

where  $\langle \cdot \rangle_{\eta^{(\pm)}}$  represents expectation values of observables with respect to the output quantum wave function generated by

the same circuit with circuit parameter  $\eta^\pm \equiv \eta \pm \frac{\pi}{2}$ . Note Eq. (A2) is an unbiased estimate of the gradient in contrast to the finite difference approaches which are sensitive to the noise of the quantum circuit.

Since the quantum circuit Born machine employs the same parametrization, we identify  $|x\rangle\langle x|$  as the observable and

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \boldsymbol{\theta}} = & \frac{1}{2} \left( \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}}(y) p_{\boldsymbol{\theta}^+}(x) - \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}}(y) p_{\boldsymbol{\theta}^-}(x) \right) + \frac{1}{2} \left( \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}}(x) p_{\boldsymbol{\theta}^+}(y) - \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}}(x) p_{\boldsymbol{\theta}^-}(y) \right) \\ & - \left( \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}^+}(x) \pi(y) - \sum_{x,y} K(x,y) p_{\boldsymbol{\theta}^-}(x) \pi(y) \right). \end{aligned} \quad (\text{A4})$$

Using the symmetric condition of the kernel  $K(x,y) = K(y,x)$ , we arrive at Eq. (2) in the main text. This gradient algorithm for the QCBM scales as  $O(d^2)$  where  $d$  is the depth of the circuit, which is less efficient compared to the linear scaling back-propagation algorithm for classical neural networks. It is still unknown whether one can reach similar scaling on a quantum computer since the back-propagation algorithm requires cached intermediate results in the forward evaluation pass. We have checked the correctness of the gradient estimator Eq. (2) against numerical finite difference.

## 2. Generalization to statistic functionals

To gain a better understanding of what kind of losses for a quantum circuit can be easily differentiated, we generalize the above result by considering an arbitrary function  $f(\mathbf{X})$  with a sequence of bit strings  $\mathbf{X} \equiv \{x_1, x_2, \dots, x_r\}$  as its arguments. Let us define the following expectation of this function:

$$\mathbb{E}_f(\boldsymbol{\Gamma}) \equiv \mathbb{E}_{\{x_i \sim p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_i}\}_{i=1}^r} [f(\mathbf{X})]. \quad (\text{A5})$$

Here,  $\boldsymbol{\Gamma} = \{\boldsymbol{\gamma}_1, \boldsymbol{\gamma}_2, \dots, \boldsymbol{\gamma}_r\}$  is the offset angles applied to circuit parameters, which means the probability distributions of generated samples is  $\{p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_1}, p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_2}, \dots, p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_r}\}$ . Writing out the above expectation explicitly, we have

$$\mathbb{E}_f(\boldsymbol{\Gamma}) = \sum_{\mathbf{X}} f(\mathbf{X}) \prod_i p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_i}(x_i), \quad (\text{A6})$$

where index  $i$  runs from 1 to  $r$ . Its partial derivative with respect to  $\theta_l^\alpha$  is

$$\frac{\partial \mathbb{E}_f(\boldsymbol{\Gamma})}{\partial \theta_l^\alpha} = \sum_{\mathbf{X}} f(\mathbf{X}) \sum_j \frac{\partial p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_j}(x_j)}{\partial \theta_l^\alpha} \prod_{i \neq j} p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_i}(x_i). \quad (\text{A7})$$

Again, using Eq. (A3), we have

$$\begin{aligned} \frac{\partial \mathbb{E}_f(\boldsymbol{\Gamma})}{\partial \theta_l^\alpha} = & \frac{1}{2} \sum_{j,s=\pm} \sum_{\mathbf{X}} f(\mathbf{X}) p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_j+s\frac{\pi}{2}\mathbf{e}_l^\alpha}(x_j) \prod_{i \neq j} p_{\boldsymbol{\theta}+\boldsymbol{\gamma}_i}(x_i) \\ = & \frac{1}{2} \sum_{j,s=\pm} \mathbb{E}_f \left( \left\{ \boldsymbol{\gamma}_i + s \delta_{ij} \frac{\pi}{2} \mathbf{e}_l^\alpha \right\}_{i=1}^r \right) \end{aligned} \quad (\text{A8})$$

apply Eq. (A2) to the output probability of the circuit,

$$\frac{\partial p_{\boldsymbol{\theta}}(x)}{\partial \theta_l^\alpha} = \frac{1}{2} [p_{\boldsymbol{\theta}^+}(x) - p_{\boldsymbol{\theta}^-}(x)], \quad (\text{A3})$$

where  $\boldsymbol{\theta}^\pm \equiv \boldsymbol{\theta} \pm \frac{\pi}{2} \mathbf{e}_l^\alpha$  with  $\mathbf{e}_l^\alpha$  the  $(l, \alpha)$ th unit vector in parameter space. Substituting Eq. (A3) into Eq. (A1) we have

If  $f$  is symmetric,  $\mathbb{E}_f(\mathbf{0})$  becomes a statistic functional [80], then Eq. (A8) can be further simplified to

$$\frac{\partial \mathbb{E}_f(\boldsymbol{\Gamma})}{\partial \theta_l^\alpha} = \frac{r}{2} \sum_{s=\pm} \mathbb{E}_f \left( \left\{ \boldsymbol{\gamma}_0 + s \frac{\pi}{2} \mathbf{e}_l^\alpha, \boldsymbol{\gamma}_1, \dots, \boldsymbol{\gamma}_r \right\} \right), \quad (\text{A9})$$

which contains only two terms. This result can be readily verified by calculating the gradient of the MMD loss, noting the expectation of a kernel function is a statistic functional of degree 2. In practice, this expectation value is estimated using the unbiased  $U$  statistics [81]. By repeatedly applying Eq. (A8), we will be able to obtain higher-order gradients.

## 3. Gradient of the KL divergence

The KL divergence [73] reads

$$\mathbb{K}\mathbb{L}(\pi \parallel p_{\boldsymbol{\theta}}) = - \sum_x \pi(x) \ln \left[ \frac{p_{\boldsymbol{\theta}}(x)}{\pi(x)} \right]. \quad (\text{A10})$$

Its gradient is

$$\frac{\partial \mathbb{K}\mathbb{L}}{\partial \theta_l^\alpha} = - \sum_x \pi(x) \frac{p_{\boldsymbol{\theta}^+}(x) - p_{\boldsymbol{\theta}^-}(x)}{2p_{\boldsymbol{\theta}}(x)}, \quad (\text{A11})$$

where the  $\frac{p_{\boldsymbol{\theta}^+}(x) - p_{\boldsymbol{\theta}^-}(x)}{p_{\boldsymbol{\theta}}(x)}$  term cannot be transformed to a sampling problem on the quantum circuit. Since minimizing the negative logarithm likelihood is equivalent to minimizing the KL divergence, it faces the same problem.

## APPENDIX B: GRADIENT ANALYSIS

In this appendix, we analyze the effect of depth and sampling error to gradients. For a nine-qubit quantum circuit of depth  $d = 100$  with random parameters used in generating the BARS-AND-STRIPES dataset in the main text. A histogram of the MMD loss gradient values in different layers are shown in Fig. 7(a). Distributions of gradients in different layers have no significant differences, thus a deep QCBM does not suffer from vanishing or exploding gradients when the layers go deeper. This property has its physical origin in how the wave function changes when a perturbation is applied to  $\theta_l^\alpha$  of the quantum circuit. Suppose we have a quantum circuit that

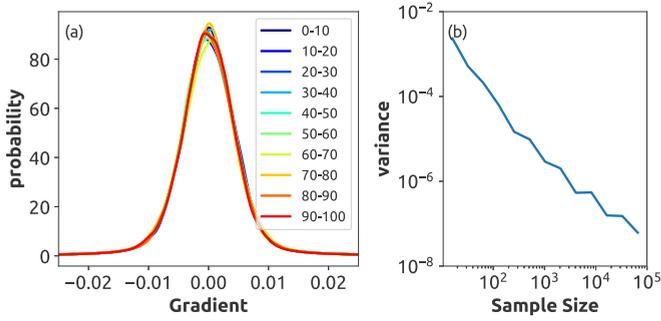


FIG. 7. (a) The probability distribution of gradient elements in different layers grouped by depth. The circuit parameters are chosen randomly. (b) The variance of the MMD loss gradients as a function of sample size  $N$ .

performing the following evolution:

$$|\psi\rangle = U_{N:1}|0\rangle. \quad (\text{B1})$$

Introduce a perturbation  $\delta$  to  $\theta_i^\alpha$ ,

$$|\psi'\rangle = U_{N:k} e^{i[(\theta_i^\alpha + \delta)/2]\Sigma_k} U_{k-1:1}|0\rangle. \quad (\text{B2})$$

Its fidelity with respect to the unperturbed wave function is

$$\begin{aligned} F &= |\langle\psi|\psi'\rangle| \\ &= \left| 1 - \frac{\delta^2}{8} + i \frac{\delta}{2} \langle\phi|\Sigma_k|\phi\rangle \right| + O(\delta^3) \\ &\simeq \sqrt{1 - \frac{\delta^2}{4}(1 - \langle\phi|\Sigma_k|\phi\rangle^2)}. \end{aligned} \quad (\text{B3})$$

Here,  $|\phi\rangle \equiv U_{k-1:1}|0\rangle$  can be assumed to be random in a deep circuit. Thus we have the fidelity susceptibility [82,83]  $\chi_F = \frac{\partial^2 \ln F}{\partial \delta^2} = 0.25(1 - \langle\Sigma_k\rangle^2)$  which is independent of  $\theta^k$ . Since we have  $\Sigma_k^2 = 1$ ,  $\chi_F$  is bounded by 0.25 and is independent of the layer index.

Next, we show the error in gradient estimation will decrease systematically as the number of samples increases for the MMD loss as can be seen from the curve of variance calculated using 100 copies of independent samples in Fig. 7(b). In our case, to give a valid estimation of gradients, the standard error should be lower than the typical amplitude of gradients. As can be referred to from Fig. 7(a), the typical amplitude of gradient is approximately  $5 \times 10^{-3}$ , thus the sample size should be larger than  $10^3$  in order to give a good estimation for these gradients.

- 
- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning* (MIT Press, Cambridge, MA, 2016).
- [2] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, [arXiv:1703.10593](https://arxiv.org/abs/1703.10593).
- [3] A. Van Den Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, [arXiv:1609.03499](https://arxiv.org/abs/1609.03499).
- [4] R. Gómez-Bombarelli, J. N. Wei, D. Duvenaud, J. M. Hernández-Lobato, B. Sánchez-Lengeling, D. Sheberla, J. Aguilera-Iparraguirre, T. D. Hirzel, R. P. Adams, and A. Aspuru-Guzik, *ACS Cent. Sci.* **4**, 268 (2018).
- [5] J. Preskill, *Quantum* **2**, 79 (2018).
- [6] W. Knight, MIT Technol. Rev. (2017), <https://www.technologyreview.com/s/609451/ibm-raises-the-bar-with-a-50-qubit-quantum-computer/>.
- [7] J. Kelly, google's new quantum processor, (2018), <https://research.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html>.
- [8] M. H. Amin, E. Andriyash, J. Rolfe, B. Kulchitsky, and R. Melko, *Phys. Rev. X* **8**, 021050 (2018).
- [9] A. Khoshaman, W. Vinci, B. Denis, E. Andriyash, and M. H. Amin, *Quantum Sci. Technol.* **4**, 014001 (2019).
- [10] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, *Phys. Rev. X* **7**, 041052 (2017).
- [11] M. Kieferová and N. Wiebe, *Phys. Rev. A* **96**, 062327 (2017).
- [12] D. H. Ackley, G. E. Hinton, and T. J. Sejnowski, *Cognitive Sci.* **9**, 147 (1985).
- [13] V. Dumoulin, I. J. Goodfellow, A. C. Courville, and Y. Bengio, *AAAI'14 Proceedings of the Twenty-Eighth AAAI Conference on Artificial Intelligence, Quebec, 2014* (AIAA, Reston, VA, 2014) Vol. 2014, pp. 1199–1205.
- [14] X. Gao, Z. Zhang, and L. Duan, [arXiv:1711.02038](https://arxiv.org/abs/1711.02038).
- [15] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, Cambridge, MA, 2009).
- [16] Z.-Y. Han, J. Wang, H. Fan, L. Wang, and P. Zhang, *Phys. Rev. X* **8**, 031012 (2018).
- [17] S. Cheng, J. Chen, and L. Wang, *Entropy* **20**, 583 (2018).
- [18] M. Benedetti, D. Garcia-Pintos, Y. Nam, A. Perdomo-Ortiz, O. Perdomo, and V. Leyton-Ortega, [arXiv:1801.07686](https://arxiv.org/abs/1801.07686).
- [19] M. Born, *Z. Phys.* **38**, 803 (1926).
- [20] L. Fortnow and J. Rogers, *J. Comput. Syst. Sci.* **59**, 240 (1999).
- [21] S. Aaronson and A. Arkhipov, *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing, San Jose, CA, 2011* (ACM, New York, 2011), pp. 333–342.
- [22] A. Lund, M. J. Bremner, and T. Ralph, *npj Quantum Inf.* **3**, 15 (2017).
- [23] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, [arXiv:1803.04402](https://arxiv.org/abs/1803.04402).
- [24] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, *Nat. Phys.* **14**, 595 (2018).
- [25] W. Huggins, P. Patel, K. B. Whaley, and E. M. Stoudenmire, [arXiv:1803.11537](https://arxiv.org/abs/1803.11537).
- [26] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2014), pp. 2672–2680.
- [27] I. Goodfellow, [arXiv:1701.00160](https://arxiv.org/abs/1701.00160).
- [28] G. E. Hinton and R. R. Salakhutdinov, *Science* **313**, 504 (2006).
- [29] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu, [arXiv:1601.06759](https://arxiv.org/abs/1601.06759).
- [30] L. Dinh, J. Sohl-Dickstein, and S. Bengio, [arXiv:1605.08803](https://arxiv.org/abs/1605.08803).
- [31] D. P. Kingma, T. Salimans, X. Chen, I. Sutskever, T. Salimans, R. Jozefowicz, and M. Welling, [arXiv:1606.04934](https://arxiv.org/abs/1606.04934).

- [32] D. J. Rezende and S. Mohamed, [arXiv:1505.05770](https://arxiv.org/abs/1505.05770).
- [33] G. Papamakarios, I. Murray, and T. Pavlakou, *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2017), pp. 2335–2344.
- [34] D. P. Kingma and M. Welling, [arXiv:1312.6114](https://arxiv.org/abs/1312.6114).
- [35] D. J. Rezende, S. Mohamed, and D. Wierstra, [arXiv:1401.4082](https://arxiv.org/abs/1401.4082).
- [36] K. Simonyan and A. Zisserman, [arXiv:1409.1556](https://arxiv.org/abs/1409.1556).
- [37] M. Minsky and S. A. Papert, *Perceptrons: An Introduction to Computational Geometry* (MIT Press, Cambridge, MA, 1969).
- [38] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, *Nature (London)* **323**, 533 (1986).
- [39] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, *Nature (London)* **549**, 195 (2017).
- [40] N. Wiebe, A. Kapoor, and K. M. Svore, [arXiv:1412.3489](https://arxiv.org/abs/1412.3489).
- [41] J. Li, X. Yang, X. Peng, and C.-P. Sun, *Phys. Rev. Lett.* **118**, 150503 (2017).
- [42] E. Farhi and H. Neven, MIT Report No. MIT-CTP/4985 (unpublished).
- [43] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, *Phys. Rev. A* **98**, 032309 (2018).
- [44] Y. Li and R. E. Turner, [arXiv:1705.07107](https://arxiv.org/abs/1705.07107).
- [45] E. Farhi, J. Goldstone, and S. Gutmann, MIT Report No. MIT-CTP/4610 (unpublished).
- [46] E. Farhi, J. Goldstone, S. Gutmann, and H. Neven, Report No. MIT-CTP/4893 (unpublished).
- [47] J. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong *et al.*, [arXiv:1712.05771](https://arxiv.org/abs/1712.05771).
- [48] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, *Nat. Commun.* **5**, 4213 (2014).
- [49] P. J. J. O’Malley, R. Babbush, I. D. Kivlichan, J. Romero, J. R. McClean, R. Barends, J. Kelly, P. Roushan, A. Tranter, N. Ding, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, P. V. Coveney, P. J. Love, H. Neven, A. Aspuru-Guzik, and J. M. Martinis, *Phys. Rev. X* **6**, 031007 (2016).
- [50] A. Kandala, A. Mezzacapo, K. Temme, M. Takita, M. Brink, J. M. Chow, and J. M. Gambetta, *Nature (London)* **549**, 242 (2017).
- [51] The number of parameters is not  $3(d+1)n$  here because the leading and trailing  $R_z$  gates can be omitted without affecting the output probability distribution.
- [52] C. Chow and C. Liu, *IEEE Trans. Inf. Theory* **14**, 462 (1968).
- [53] G. H. Low, T. J. Yoder, and I. L. Chuang, *Phys. Rev. A* **89**, 062315 (2014).
- [54] A. Zulehner, A. Paler, and R. Wille, [arXiv:1712.04722](https://arxiv.org/abs/1712.04722).
- [55] S. Mohamed and B. Lakshminarayanan, [arXiv:1610.03483](https://arxiv.org/abs/1610.03483).
- [56] N. Anderson, P. Hall, and D. Titterton, *J. Multivariate Anal.* **50**, 41 (1994).
- [57] A. Gretton, K. M. Borgwardt, M. Rasch, B. Schölkopf, and A. J. Smola, *Advances in Neural Information Processing Systems* (MIT Press, Cambridge, MA, 2007), pp. 513–520.
- [58] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola, *JMLR* **13**, 723 (2012).
- [59] T. Hofmann, B. Schölkopf, and A. J. Smola, *Ann. Statist.* **36**, 1171 (2008).
- [60] Y. Li, K. Swersky, and R. Zemel, [arXiv:1502.02761](https://arxiv.org/abs/1502.02761).
- [61] G. K. Dziugaite, D. M. Roy, and Z. Ghahramani, [arXiv:1505.03906](https://arxiv.org/abs/1505.03906).
- [62] C.-L. Li, W.-C. Chang, Y. Cheng, Y. Yang, and B. Póczos, [arXiv:1705.08584](https://arxiv.org/abs/1705.08584).
- [63] J. C. Spall, *IEEE Trans. Aerosp. Electron. Syst.* **34**, 817 (1998).
- [64] For the small dataset employed in our numerical experiment, we can afford to average over the whole training dataset.
- [65] <https://github.com/GiggleLiu/QuantumCircuitBornMachine>.
- [66] <https://quantumbfs.github.io/Yao.jl/latest/tutorial/QCBM/>.
- [67] D. J. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge University Press, Cambridge, UK, 2003).
- [68] R. H. Byrd, P. Lu, J. Nocedal, and C. Zhu, *SIAM J. Sci. Comput.* **16**, 1190 (1995).
- [69] D. P. Kingma and J. Ba, [arXiv:1412.6980](https://arxiv.org/abs/1412.6980).
- [70] N. Hansen, in *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, edited by J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea (Springer, Berlin/Heidelberg, 2006), pp. 75–102.
- [71] L. M. Rios and N. V. Sahinidis, *J. Global Optim.* **56**, 1247 (2013).
- [72] Note for this circuit of depth  $d = 10$  that has 279 parameters, we need to make  $279 \times N \times 2 + 1$  measurements in a gradient estimation.
- [73] S. Kullback and R. A. Leibler, *Ann. Math. Stat.* **22**, 79 (1951).
- [74] A. Choromanska, M. Henaff, M. Mathieu, G. B. Arous, and Y. LeCun, *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics, San Diego, CA, 2015* (MIR, Moscow, 2015), pp. 192–204.
- [75] K. He, X. Zhang, S. Ren, and J. Sun, *IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Las Vegas, 2016* (IEEE, Piscataway, NJ, 2016), pp. 770–778.
- [76] L. Jing, Y. Shen, T. Dubcek, J. Peurifoy, S. Skirlo, Y. LeCun, M. Tegmark, and M. Soljagic, [arXiv:1612.05231](https://arxiv.org/abs/1612.05231).
- [77] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, [arXiv:1803.11173](https://arxiv.org/abs/1803.11173).
- [78] D. S. Steiger, T. Häner, and M. Troyer, *Quantum* **2**, 49 (2018).
- [79] T. Häner, D. S. Steiger, K. Svore, and M. Troyer, *Quantum Sci. Technol.* **3**, 020501 (2018).
- [80] R. v. Mises, *Ann. Math. Statist.* **18**, 309 (1947).
- [81] W. Hoeffding, *Breakthroughs in Statistics* (Springer, Berlin, 1992), pp. 308–334.
- [82] W.-L. You, Y.-W. Li, and S.-J. Gu, *Phys. Rev. E* **76**, 022101 (2007).
- [83] L. Wang, Y.-H. Liu, J. Imriška, P. N. Ma, and M. Troyer, *Phys. Rev. X* **5**, 031007 (2015).