Variational Inference with a Quantum Computer

Marcello Benedetti^{1,*} Brian Coyle^{1,2} Mattia Fiorentini¹
Michael Lubasch¹ and Matthias Rosenkranz^{1,†}

¹*Cambridge Quantum Computing Limited, London SW1E 6DR, United Kingdom*

²*School of Informatics, University of Edinburgh, Edinburgh EH8 9AB, United Kingdom*

(Received 26 April 2021; revised 17 September 2021; accepted 28 September 2021; published 28 October 2021)

Inference is the task of drawing conclusions about unobserved variables given observations of related variables. Applications range from identifying diseases from symptoms to classifying economic regimes from price movements. Unfortunately, performing exact inference is intractable in general. One alternative is variational inference, where a candidate probability distribution is optimized to approximate the posterior distribution over unobserved variables. For good approximations, a flexible and highly expressive candidate distribution is desirable. In this work, we use quantum Born machines as variational distributions over discrete variables. We apply the framework of operator variational inference to achieve this goal. In particular, we adopt two specific realizations: one with an adversarial objective and one based on the kernelized Stein discrepancy. We demonstrate the approach numerically using examples of Bayesian networks, and implement an experiment on an IBM quantum computer. Our techniques enable efficient variational inference with distributions beyond those that are efficiently representable on a classical computer.

DOI: [10.1103/PhysRevApplied.16.044057](https://doi.org/10.1103/PhysRevApplied.16.044057)

I. INTRODUCTION

Probabilistic graphical models describe the dependencies of random variables in complex systems [1]. This framework enables two important tasks: learning and inference. Learning yields a model that approximates the observed data distribution. Inference uses the model to answer queries about unobserved variables given observations of other variables. In general, exact inference is intractable and so producing good *approximate* solutions becomes a desirable goal. This article introduces approximate inference solutions using a hybrid quantum-classical framework.

Prominent examples of probabilistic graphical models are Bayesian and Markov networks. Applications across many domains employ inference on those models, including in health care and medicine [2,3], biology, genetics, and forensics [4,5], finance [6], and fault diagnosis [7].

These are applications where qualifying and quantifying the uncertainty of conclusions is crucial. The posterior distribution can be used to quantify this uncertainty. It can also be used in other downstream tasks such as determining the likeliest configuration of unobserved variables that best explains observed data.

Approximate inference methods broadly fall into two categories: Markov chain Monte Carlo (MCMC) and variational inference (VI). MCMC methods produce samples from the true posterior distribution in an asymptotic limit [8,9]. VI is a machine learning technique that casts inference as an optimization problem over a parameterized family of probability distributions [10]. If quantum computers can deliver even a small improvement to these methods, the impact across science and engineering could be large.

Initial progress in combining MCMC with quantum computing was achieved by replacing standard MCMC with quantum annealing hardware in the training of some types of Bayesian and Markov networks [11–17]. Despite promising empirical results, it proved difficult to show *if* and *when* a quantum advantage could be delivered. An arguably more promising path towards quantum advantage is through gate-based quantum computers. In this context, there exist algorithms for MCMC with proven asymptotic advantage [18–21], but they require error correction and other features that go beyond the capability of existing machines. Researchers have recently shifted their focus to

*marcello.benedetti@cambridgequantum.com

†matthias.rosenkranz@cambridgequantum.com

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

near-term machines, proposing new quantum algorithms for sampling thermal distributions [22–26].

The advantage of combining quantum computing with VI has not been explored to the same extent as MCMC. Previous work includes classical VI algorithms using ideas from quantum annealing [27–29]. In this work, we turn our attention to performing VI on a quantum computer. We adopt a distinct approach that focuses on improving inference in classical probabilistic models by using quantum resources. We use *Born machines*, which are quantum machine learning models that exhibit high expressivity. We show how to employ gradient-based methods and *amortization* in the training phase. These choices are inspired by recent advances in classical VI [30].

Finding a quantum advantage in machine learning in *any* capacity is an exciting research goal, and promising theoretical works in this direction have recently been developed across several prominent subfields. Advantages in supervised learning have been proposed considering: information theoretic arguments [31–34], probably approximately correct (PAC) learning [35–38], and the representation of data in such models [39–41]. More relevant for our purposes are results in *unsupervised* learning, which have considered complexity and learning theory arguments for distributions [42–44] and quantum nonlocality and contextuality [45]. Furthermore, some advantages have been observed experimentally [46–49]. For further reading, see Refs. [50–52] for recent overviews of some advances in quantum machine learning.

In Sec. II we describe VI and its applications. In Sec. III we describe using Born machines to approximate posterior distributions. In Sec. IV we use the framework of operator VI to derive two suitable objective functions, and we employ classical techniques to deal with the problematic terms in the objectives. In Sec. V we demonstrate the methods on Bayesian networks. We conclude in Sec. VI with a discussion of possible generalizations and future work.

II. VARIATIONAL INFERENCE AND APPLICATIONS

It is important to clarify what type of inference we are referring to. Consider a probabilistic model p over some set of random variables, $\mathcal{Y}$. The variables in $\mathcal{Y}$ can be continuous or discrete. Furthermore, assume that we are given *evidence* for some variables in the model. This set of variables, denoted $\mathcal{X} \subseteq \mathcal{Y}$, is then *observed* (fixed to the values of the evidence) and we use the vector notation $\mathbf{x}$ to denote a realization of these observed variables. We now want to *infer* the posterior distribution of the unobserved variables, those in the set $\mathcal{Z} := \mathcal{Y} \setminus \mathcal{X}$. Denoting these by a vector $\mathbf{z}$, our target is that of computing the posterior distribution $p(\mathbf{z}|\mathbf{x})$, the conditional probability of $\mathbf{z}$ given $\mathbf{x}$. By definition, the conditional can be expressed in terms of the joint divided by the marginal: $p(\mathbf{z}|\mathbf{x}) = p(\mathbf{x}, \mathbf{z})/p(\mathbf{x})$. Also,

recall that the joint can be written as $p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$. Bayes’ theorem combines the two identities and yields $p(\mathbf{z}|\mathbf{x}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})/p(\mathbf{x})$.

As described above, the inference problem is rather general. To set the scene, let us discuss some concrete examples in the context of Bayesian networks. These are commonly targeted in inference problems, and as such shall be our test bed for this work. A Bayesian network describes a set of random variables with a clear conditional probability structure. This structure is a directed acyclic graph where the conditional probabilities are modeled by tables, by explicit distributions, or even by neural networks. Figure 1(a) is the textbook example of a Bayesian network for the distribution of binary variables: cloudy (C), sprinkler (S), rain (R), and the grass being wet (W). According to the graph this distribution factorizes as $P(C, S, R, W) = P(C)P(S|C)P(R|C)P(W|S, R)$. A possible

FIG. 1. Some applications of inference on Bayesian networks. Filled (empty) circles indicate observed (unobserved) random variables. (a) “Sprinkler” network. (b) Regime switching in time series. (c) “Lung cancer” network [53]. (d) Error correction codes.

inferential question is: what is the probability distribution of C , S , and R given that $W = \text{tr}$? This can be estimated by “inverting” the probabilities using Bayes’ theorem: $p(C, S, R|W = \text{tr}) = p(W = \text{tr}|C, S, R)p(C, S, R)/p(W)$.

Figures 1(b)–1(d) show a few additional applications of inference on Bayesian networks. The hidden Markov model in Fig. 1(b) describes the joint probability distribution of a time series of asset returns and an unobserved “market regime” (e.g., a booming versus a recessive economic regime). A typical inference task is to detect regime switches by observing asset returns [54]. Figure 1(c) illustrates a modified version of the “lung cancer” Bayesian network that is an example from medical diagnosis (see, e.g., Ref. [3] and the references therein). This network encodes expert knowledge about the relationship between risk factors, diseases, and symptoms. In health care, careful design of the network and algorithm are critical in order to reduce biases, e.g., in relation to health care access [55]. Note that inference in medical diagnosis is often causal instead of associative [3]. Bayesian networks can be interpreted causally and help answer causal queries using Pearl’s do-calculus [56]. Finally, Fig. 1(d) shows a Bayesian network representation of turbo codes, which are error correction schemes used in 3G and 4G mobile communications. The inference task for the receiver is to recover the original information bits from the information bits and codewords received over a noisy channel [57].

Inference is a computationally hard task in all but the simplest probabilistic models. Roth [58] extended results of Cooper [59] and showed that *exact* inference in Bayesian networks with discrete variables is sharp-P-complete. Dagum and Luby [60] showed that even *approximate* inference is NP-hard. Therefore, unless some particular constraints are in place, these calculations are intractable. In many cases one is able to perform a “forward pass” and obtain unbiased samples from the joint $(\mathbf{x}, \mathbf{z}) \sim p(\mathbf{x}, \mathbf{z}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})$. However, obtaining unbiased samples from the posterior $\mathbf{z} \sim p(\mathbf{z}|\mathbf{x}) = p(\mathbf{x}, \mathbf{z})/p(\mathbf{x})$ is intractable due to the unknown normalization constant. One can perform MCMC sampling by constructing an ergodic Markov chain whose stationary distribution is the desired posterior. MCMC methods have nice theoretical guarantees, but they may converge slowly in practice [9]. In contrast, VI is often faster in high dimensions but does not come with guarantees. The idea in VI is to optimize a variational distribution q by minimizing its “distance” from the true posterior p (see Ref. [10] for an introduction to the topic and Ref. [30] for a review of the recent advances).

VI has experienced a resurgence in recent years due to substantial developments. First, generic methods have reduced the amount of analytical calculations required and have made VI much more user friendly (e.g., black-box VI [61]). Second, machine learning has enabled the use of highly expressive variational distributions implemented

via neural networks [62], probabilistic programs [63], nonparametric models [64], normalizing flows [65], and others. In contrast, the original VI methods were mostly limited to analytically tractable distributions such as those that could be factorized. Third, amortization methods have reduced the costs by optimizing $q_\theta(\mathbf{z}|\mathbf{x})$ where the vector of parameters θ is “shared” among all possible observations $\mathbf{x}$ instead of optimizing individual parameters for each observation. This approach can also generalize across inferences.

Casting the inference problem as an optimization problem comes itself with challenges, in particular when the unobserved variables are discrete. REINFORCE [66] is a generic method that requires calculation of the score $\partial_\theta \log q_\theta(\mathbf{z}|\mathbf{x})$ and may suffer from high variance. Gumbel-softmax reparameterizations [67,68] use a continuous relaxation that does not follow the exact distribution.

We now show that near-term quantum computers provide an alternative tool for VI. We use the quantum Born machine as a candidate for variational distributions. These models are highly expressive and they naturally represent discrete distributions as a result of quantum measurement. Furthermore, the Born machine can be trained by gradient-based optimization.

III. BORN MACHINES AS IMPLICIT VARIATIONAL DISTRIBUTIONS

By exploiting the inherent probabilistic nature of quantum mechanics, one can model the probability distribution of classical data using a pure quantum state. This model based on the Born rule in quantum mechanics is referred to as the Born machine [69]. Let us consider binary vectors $\mathbf{z} \in \{0, 1\}^n$, where n is the number of variables. The Born machine is a normalized quantum state $|\psi(\theta)\rangle$ parameterized by θ that outputs n -bit strings with probabilities $q_\theta(\mathbf{z}) = |\langle \mathbf{z} | \psi(\theta) \rangle|^2$. Here $|\mathbf{z}\rangle$ are computational basis states, thus sampling the above probability boils down to a simple measurement. Other forms of discrete variables can be dealt with by a suitable encoding. When using amortization, the variational distribution requires conditioning on observed variables. To extend the Born machine and include this feature, we let $\mathbf{x}$ play the role of additional parameters. This yields a pure state where the output probabilities are $q_\theta(\mathbf{z}|\mathbf{x}) = |\langle \mathbf{z} | \psi(\theta, \mathbf{x}) \rangle|^2$. Figure 2 shows the relation between the classical model and the quantum model for approximate inference.

Born machines have been applied for benchmarking hybrid quantum-classical systems [70–73], generative modeling [74–76], finance [47,48,77], anomaly detection [78], and have been proposed for demonstrating quantum advantage [43,79]. These models can be realized in a variety of ways, in both classical and quantum computers. When realized via certain classes of quantum circuits, they are classically intractable to simulate [80]. For example,

FIG. 2. Probabilistic models used in our methods. The classical model comprises a prior over unobserved discrete variables and a likelihood over observed variables. The quantum model approximates the posterior distribution of the unobserved variables given observed data. All distributions can be sampled efficiently as long as one follows the arrows and uses a suitable computer.

instantaneous quantum polytime (IQP) circuits are Born machines with $O[\text{poly}(n)]$ parameters that yield classically intractable distributions in the average case under widely accepted complexity theoretic assumptions [81]. Additional examples for such classically hard circuits are boson sampling [82] and random circuits [83,84]. Thus, quantum Born machines have expressive power larger than that of classical models, including neural networks [85] and partially matrix product states [86]. It can be shown that the model remains classically intractable throughout training [43]. We return to this discussion in Appendix A.

A useful way to classify probabilistic models is the following: *prescribed* models provide an explicit parametric specification of the distribution, *implicit* models define only the data generation process [87]. Born machines can be effectively regarded as implicit models. It is easy to obtain an unbiased sample as long as we can execute the corresponding circuit on a quantum computer and measure the computational basis. However, it requires exponential resources to estimate the probability of a sample with multiplicative error.

Implicit models are challenging to train with standard methods. The major challenge is the design of the objective function precisely because likelihoods are “prohibited.” Valid objectives involve only statistical quantities (such as expectation values) that can be efficiently estimated from samples [90]. For generative modeling with Born machines, progress has been made towards practical objectives such as moment matching [70], maximum mean discrepancy [91], Stein and Sinkhorn divergences [43], adversarial objectives [92,93], as well as incorporating Bayesian priors on the parameters [85].

In the next section we make some progress towards practical objectives for VI with Born machines. First, we mention some related work. The approaches in Ref. [23] use VI ideas to deal with thermal states and quantum data. In contrast, our inference methods apply to classical graphical models and classical data. The approaches of Refs. [18,94,95] aim to encode Bayesian networks

directly in a quantum state, and subsequently perform exact inference. In contrast, our inference methods are approximate, efficient, and apply also to graphical models that are not Bayesian networks.

IV. OPERATOR VARIATIONAL INFERENCE

Operator variational inference (OPVI) [96] is a rather general method that uses mathematical operators to design objectives for the approximate posterior. Suitable operators are those for which (i) the minima of the variational objective is attained at the true posterior, and (ii) it is possible to estimate the objective without computing the true posterior. In general, the amortized OPVI objective is

$$\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \sup_{f \in \mathcal{F}} h\{\mathbb{E}_{\mathbf{z} \sim q(\mathbf{z}|\mathbf{x})}[(O^{p,q}f)(\mathbf{z})]\}, \quad (1)$$

where $f(\cdot) \in \mathbb{R}^d$ is a test function within the chosen family $\mathcal{F}$, $O^{p,q}$ is an operator that depends on $p(\mathbf{z}|\mathbf{x})$ and $q(\mathbf{z}|\mathbf{x})$, and $h(\cdot) \in [0, \infty]$ yields a non-negative objective. Note that we have an expectation over the data distribution $p_{\mathcal{D}}(\mathbf{x})$. This indicates the average over a dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_i$ of observations to condition on.

We present two methods that follow directly from two operator choices. These operator choices result in objectives based on the Kullback-Leibler (KL) divergence and the Stein discrepancy. The former utilizes an *adversary* to make the computation tractable, whereas in the latter, the tractability arises from the use of a *kernel* function.

The KL divergence is an example of an *f* divergence, while the Stein discrepancy is in the class of *integral probability metrics*, two fundamentally different families of probability distance measures. OPVI can yield methods from these two different classes under a suitable choice of operator. As shown in Ref. [97], these two families intersect nontrivially only at the total variation distance (TVD). It is for this reason that we choose the TVD as our benchmark in later numerical results. Integral probability metrics [43,91] and *f* divergences [70,73] have both been used to train Born machines for the task of generative modeling, and we allow the more thorough incorporation of these methods in future work. Furthermore, these works demonstrate that the choice of training metric has a significant effect on the trainability of the model. It is for this reason that we adopt the OPVI framework in generality, rather than focusing on any specific instance of it. This allows the switching between different metrics depending on the need.

A. The adversarial method

One possible objective function for VI is the KL divergence of the true posterior relative to the approximate one. This is obtained from Eq. (1) by choosing h to be the identity function, and by choosing the operator $(O^{p,q}f)(\mathbf{z}) =$

$\log[q(\mathbf{z}|\mathbf{x})/p(\mathbf{z}|\mathbf{x})]$ for all f :

$$\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \underbrace{\mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} \left[\log \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{p(\mathbf{z}|\mathbf{x})} \right]}_{\text{Kullback-Leibler divergence}}. \quad (2)$$

Here q_{θ} is the variational distribution parameterized by θ . To avoid singularities we assume that $p(\mathbf{z}|\mathbf{x}) > 0$ for all $\mathbf{x}$ and all $\mathbf{z}$. The objective's minimum is zero and is attained by the true posterior $q_{\theta}(\mathbf{z}|\mathbf{x}) = p(\mathbf{z}|\mathbf{x})$ for all $\mathbf{x}$.

There are many ways to rewrite this objective. Here we focus on the form known as *prior contrastive*. Substituting Bayes' formula $p(\mathbf{z}|\mathbf{x}) = p(\mathbf{x}|\mathbf{z})p(\mathbf{z})/p(\mathbf{x})$ into the equation above we obtain

$$\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} \left[\log \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} - \log p(\mathbf{x}|\mathbf{z}) \right] + \text{const.}, \quad (3)$$

where we ignore the term $\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} [\log p(\mathbf{x})]$ since it is constant with respect to q . This objective has been used in many generative models, including the celebrated variational autoencoder (VAE) [62,98]. While the original VAE relies on tractable variational posteriors, one could also use more powerful, implicit distributions as demonstrated in adversarial variational Bayes [99] and in prior-contrastive adversarial VI [100].

Let us use a Born machine to model the variational distribution $q_{\theta}(\mathbf{z}|\mathbf{x}) = |\langle \mathbf{z} | \psi(\theta, \mathbf{x}) \rangle|^2$. Since this model is implicit, the ratio $q_{\theta}(\mathbf{z}|\mathbf{x})/p(\mathbf{z})$ in Eq. (3) cannot be computed efficiently. Therefore, we introduce an adversarial method for approximately estimating the ratio above. The key insight here is that this ratio can be estimated from the output of a binary classifier [87]. Suppose that we ascribe samples $(\mathbf{z}, \mathbf{x}) \sim q_{\theta}(\mathbf{z}|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})$ to the first class, and samples $(\mathbf{z}, \mathbf{x}) \sim p(\mathbf{z})p_{\mathcal{D}}(\mathbf{x})$ to the second class. A binary classifier d_{ϕ} parameterized by ϕ outputs the probability $d_{\phi}(\mathbf{z}, \mathbf{x})$ that the pair $(\mathbf{z}, \mathbf{x})$ belongs to one of two classes. Hence, $1 - d_{\phi}(\mathbf{z}, \mathbf{x})$ indicates the probability that $(\mathbf{z}, \mathbf{x})$ belongs to the other class. There exist many possible choices of objective function for the classifier [87]. In this work we consider the cross-entropy

$$\mathcal{G}_{\text{KL}}(\phi; \theta) = \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\log d_{\phi}(\mathbf{z}, \mathbf{x})] + \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z})} \{\log[1 - d_{\phi}(\mathbf{z}, \mathbf{x})]\}. \quad (4)$$

The optimal classifier that maximizes this equation is [99]

$$d^*(\mathbf{z}, \mathbf{x}) = \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{q_{\theta}(\mathbf{z}|\mathbf{x}) + p(\mathbf{z})}. \quad (5)$$

Since the probabilities in Eq. (5) are unknown, the classifier must be trained on a dataset of samples. This does not pose a problem because samples from the Born machine

$q_{\theta}(\mathbf{z}|\mathbf{x})$ and prior $p(\mathbf{z})$ are easy to obtain by assumption. Once the classifier is trained, the logit transformation provides the log odds of a data point coming from the Born machine joint $q_{\theta}(\mathbf{z}|\mathbf{x})p_{\mathcal{D}}(\mathbf{x})$ versus the prior joint $p(\mathbf{z})p_{\mathcal{D}}(\mathbf{x})$. The log odds are an approximation to the log ratio of the two distributions, i.e.,

$$\text{logit}[d_{\phi}(\mathbf{z}, \mathbf{x})] \equiv \log \frac{d_{\phi}(\mathbf{z}, \mathbf{x})}{1 - d_{\phi}(\mathbf{z}, \mathbf{x})} \approx \log \frac{q_{\theta}(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})}, \quad (6)$$

which is exact if d_{ϕ} is the optimal classifier in Eq. (5). Now, we can avoid the computation of the problematic term in the KL divergence. Substituting this result into Eq. (3) and ignoring the constant term, the final objective for the Born machine is

$$\mathcal{L}_{\text{KL}}(\theta; \phi) = \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} \{\text{logit}[d_{\phi}(\mathbf{z}, \mathbf{x})] - \log p(\mathbf{x}|\mathbf{z})\}. \quad (7)$$

The optimization can be performed in tandem as

$$\begin{aligned} & \max_{\phi} \mathcal{G}_{\text{KL}}(\phi; \theta), \\ & \min_{\theta} \mathcal{L}_{\text{KL}}(\theta; \phi), \end{aligned} \quad (8)$$

using gradient ascent and descent, respectively. It can be shown [99,100] that the gradient of $\log[q_{\theta}(\mathbf{z}|\mathbf{x})/p(\mathbf{z})]$ with respect to θ vanishes. Thus, under the assumption of an optimal classifier d_{ϕ} , the gradient of Eq. (7) is significantly simplified. This gradient is derived in Appendix B.

A more intuitive interpretation of the procedure just described is as follows. The log likelihood in Eq. (3) can be expanded as $\log p(\mathbf{x}|\mathbf{z}) = \log[p(\mathbf{z}|\mathbf{x})/p(\mathbf{z})] + \log p(\mathbf{x})$. Then Eq. (3) can be rewritten as $\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim q_{\theta}(\mathbf{z}|\mathbf{x})} \{\log[q_{\theta}(\mathbf{z}|\mathbf{x})/p(\mathbf{z})] - \log[p(\mathbf{z}|\mathbf{x})/p(\mathbf{z})]\}$, revealing the difference between two log odds. The first is given by the optimal classifier in Eq. (6) for the approximate posterior and prior. The second is given by a hypothetical classifier for the true posterior and prior. The adversarial method is illustrated in Fig. 3.

B. The kernelized method

Another possible objective function for VI is the Stein discrepancy (SD) of the true posterior from the approximate one. This is obtained from Eq. (1), assuming that the image of f has the same dimension as $\mathbf{z}$, choosing h to be the absolute value, and choosing $O^{p,q}$ to be a Stein operator. A Stein operator is independent from q and is characterized by having zero expectation under the true posterior for all functions f in the chosen family $\mathcal{F}$.

FIG. 3. Adversarial variational inference with a Born machine. Step 1 optimizes a classifier d_ϕ to output the probabilities that the observed samples come from the Born machine rather than from the prior. Step 2 optimizes the Born machine q_θ to better match the true posterior. The updated Born machine is fed back into step 1 and the process repeats until convergence of the Born machine.

For binary variables, a possible Stein operator is $(O^p f)(\mathbf{z}) = s_p(\mathbf{x}, \mathbf{z})^T f(\mathbf{z}) - \text{tr}[\Delta f(\mathbf{z})]$, where

$$[s_p(\mathbf{x}, \mathbf{z})]_i = \frac{\Delta_{z_i} p(\mathbf{x}, \mathbf{z})}{p(\mathbf{x}, \mathbf{z})} = 1 - \frac{p(\mathbf{x}, \neg_i \mathbf{z})}{p(\mathbf{x}, \mathbf{z})} \quad (9)$$

is the difference score function. The partial difference operator is defined for any scalar function g on binary vectors as

$$\Delta_{z_i} g(\mathbf{z}) = g(\mathbf{z}) - g(\neg_i \mathbf{z}), \quad (10)$$

where $\neg_i \mathbf{z}$ flips the i th bit in binary vector $\mathbf{z}$. We have also defined $\Delta f(\mathbf{z})$ as the matrix with entries $[\Delta f(\mathbf{z})]_{ij} = \Delta_{z_i} f_j(\mathbf{z})$. Under the assumption that $p(\mathbf{z}|\mathbf{x}) > 0$ for all $\mathbf{x}$ and all $\mathbf{z}$, we show in Appendix C that this is a valid Stein operator for binary variables. For more general discrete variables, we refer the interested reader to Ref. [101]. Substituting these definitions into Eq. (1) we obtain [102]:

$$\mathbb{E}_{\mathbf{x} \sim p_D(\mathbf{x})} \sup_{f \in \mathcal{F}} |\underbrace{\mathbb{E}_{\mathbf{z} \sim q_\theta(\mathbf{z}|\mathbf{x})} \{s_p(\mathbf{x}, \mathbf{z})^T f(\mathbf{z}) - \text{tr}[\Delta f(\mathbf{z})]\}}_{\text{Stein discrepancy}}|. \quad (11)$$

At this point one can parameterize the test function f and obtain an adversarial objective similar in spirit to that presented in Sec. IV A. Here, however, we take a different route. If we take $\mathcal{F}$ to be a reproducing kernel Hilbert space of vector-valued functions, $\mathcal{H}$, and restricting the Hilbert space norm of f to be at most one ($\|f\|_{\mathcal{H}} \leq 1$), the supremum in Eq. (11) can be calculated in closed form using a kernel [101]. This is achieved using the “reproducing” properties of kernels, $f(\mathbf{z}) = \langle k(\mathbf{z}, \cdot), f(\cdot) \rangle_{\mathcal{H}}$, substituting into Eq. (11) and explicitly solving for the supremum (the

explicit form of this supremal f^* is not illuminating for our discussion; see [101] for details). Inserting this f^* into the SD results in the kernelized Stein discrepancy (KSD),

$$\mathbb{E}_{\mathbf{x} \sim p_D(\mathbf{x})} \underbrace{\sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_\theta(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]}}_{\text{kernelized Stein discrepancy}} \quad (12)$$

where κ_p is a so-called Stein kernel. For binary variables, this can be written as

$$\begin{aligned} \kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x}) &= s_p(\mathbf{x}, \mathbf{z})^T k(\mathbf{z}, \mathbf{z}') s_p(\mathbf{x}, \mathbf{z}') \\ &\quad - s_p(\mathbf{x}, \mathbf{z})^T \Delta_{\mathbf{z}'} k(\mathbf{z}, \mathbf{z}') - \Delta_{\mathbf{z}} k(\mathbf{z}, \mathbf{z}')^T s_p(\mathbf{x}, \mathbf{z}') \\ &\quad + \text{tr}[\Delta_{\mathbf{z}, \mathbf{z}'} k(\mathbf{z}, \mathbf{z}')]. \end{aligned} \quad (13)$$

Here $\Delta_{\mathbf{z}} k(\mathbf{z}, \cdot)$ is the vector of partial differences with components $\Delta_{z_i} k(\mathbf{z}, \cdot)$, $\Delta_{\mathbf{z}, \mathbf{z}'} k(\mathbf{z}, \mathbf{z}')$ is the matrix produced by applying $\Delta_{\mathbf{z}}$ to each component of the vector $\Delta_{\mathbf{z}'} k(\cdot, \mathbf{z}')$, and $\text{tr}(\cdot)$ is the usual trace operation on this matrix. The expression $\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})$ results in a scalar value.

Equation (12) shows that the discrepancy between the approximate and true posteriors can be calculated from samples of the approximate posterior alone. Note that the Stein kernel, κ_p , depends on *another* kernel k . For n unobserved Bernoulli variables, one possibility is the generic Hamming kernel $k(\mathbf{z}, \mathbf{z}') = \exp[-(1/n)\|\mathbf{z} - \mathbf{z}'\|_1]$. The KSD is a valid discrepancy measure if this “internal” kernel, k , is positive definite, which is the case for the Hamming kernel [101].

FIG. 4. Kernelized Stein variational inference with a Born machine. Step 1 computes the Stein discrepancy (SD) between the Born machine and the true posterior using samples from the Born machine alone. By choosing a reproducing kernel Hilbert space $\mathcal{H}$ with kernel k , the kernelized Stein discrepancy (KSD) can be calculated in closed form. Step 2 optimizes the Born machine to reduce the discrepancy. The process repeats until convergence.

In summary, constraining $\|f\|_{\mathcal{H}} \leq 1$ in Eq. (11) and substituting the KSD we obtain

$$\mathcal{L}_{\text{KSD}}(\theta) = \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]}, \quad (14)$$

and the problem consists of finding $\min_{\theta} \mathcal{L}_{\text{KSD}}(\theta)$. The gradient of Eq. (14) is derived in Appendix D. The kernelized method is illustrated in Fig. 4.

The KSD was used in Ref. [43] for generative modeling with Born machines. In that context the distribution p is unknown; thus, the authors derived methods to approximate the score s_p from available data. VI is a suitable application of the KSD because in this context we do know the joint $p(\mathbf{z}, \mathbf{x})$. Moreover the score in Eq. (9) can be computed efficiently even if we have the joint in unnormalized form.

V. EXPERIMENTS

To demonstrate our approach, we employ three experiments as the subject of the following sections. First, we validate both methods using the canonical “sprinkler” Bayesian network. Second, we consider continuous observed variables in a hidden Markov model (HMM), and illustrate how multiple observations can be incorporated effectively via amortization. Finally, we consider a larger model, the “lung cancer” Bayesian network, and demonstrate our methods using an IBM quantum computer.

A. Proof of principle with the “sprinkler” Bayesian network

As a first experiment, we classically simulate the methods on the “sprinkler” network in Fig. 1(a), one of the simplest possible examples. The purpose is to show that both methods are expected to work well even when using limited resources and without much fine tuning. First, we randomly generate the entries of each probability table from the uniform distribution $\mathcal{U}([0.01, 0.99])$ and produce a total of 30 instances of this network. For each instance, we condition on “Grass wet” being true, which means that our data distribution is $p_{\mathcal{D}}(W) = \delta(W = \text{tr})$. We infer the posterior of the remaining three variables using the three-qubit version of the hardware-efficient ansatz shown in Fig. 5. We use a layer of Hadamard gates as a state preparation, $S(\mathbf{x}) = H \otimes H \otimes H$, and initialize all parameters to approximately 0. Such hardware-efficient *ansätze*, while simple, have been shown to be vulnerable to *barren plateaus* [103–106], or regions of exponentially vanishing gradient magnitudes that make training untenable. Alternatively, one could consider *ansätze* that have been shown to be somewhat “immune” to this phenomenon [107, 108], but we leave such investigation to future work.

For the KL objective, we utilize a multilayer perceptron (MLP) classifier made of three input units, six hidden

FIG. 5. Example of hardware-efficient ansatz for the Born machine used in the experiments. All rotations R_x , R_z are parameterized by individual angles. The layer denoted as $S(\mathbf{x})$ encodes the values of the observed variables $\mathbf{x}$. The particular choice of $S(\mathbf{x})$ depends on the application (see the main text). The corresponding classical random variables are indicated below the circuit.

rectified linear units, and one sigmoid output unit. The classifier is trained with a dataset of 100 samples from the prior $p(C, R, S)$ and 100 samples from the Born machine $q(C, R, S|W = \text{tr})$. Here we use stochastic gradient descent with batches of size ten and a learning rate of 0.03. For the KSD objective, we use the Hamming kernel as above. For both methods, the Born machine is trained using 100 samples to estimate each expectation value for the gradients, and using vanilla gradient descent with a small learning rate of 0.003. We compute the TVD of the true and approximate posteriors at each epoch. We emphasize that TVD cannot be efficiently computed in general and can be shown only for small examples. Figures 6(a) [Fig. 6(b)] shows the median TVD out of the 30 instances for 1000 epochs of training for the KL (KSD) objective. A shaded area indicates a 68% confidence interval obtained by bootstrapping. Increasing the number of layers leads to better approximations to the posterior in all cases.

We compare the representation power of Born machines to that of fully factorized posteriors (e.g., those used in the naïve mean-field approximation). We perform an exhaustive search for the best full factorization $q(\mathbf{z}) = \prod_{k=1}^3 q_k(z_k)$. This is computationally intractable, but can be done for our small examples. The horizontal gray lines in Figs. 6(a) and 6(b) show the median TVD across the 30 instances for the full factorization. The 0-layer Born machine generates unentangled states and could in principle match the horizontal line. This does not happen, which indicates the challenges of optimizing a fully factorized Born machine. On the other hand, Born machines with $L > 0$ layers trained with the KL objective systematically outperform the full factorization. Born machines trained with

FIG. 6. Total variation distance of the approximate and true posteriors. Median (lines) and 68% confidence interval (shaded area) of 30 instances for (a) KL and (b) KSD objective functions for a different number of layers in the ansatz (cf. Fig. 5). Each instance is a “sprinkler” Bayesian network as in Fig. 1(a) where the conditional probabilities are chosen at random. (a) Adversarial objective. (b) Kernelized Stein objective.

the KSD objective yield suboptimal results in comparison to the KL objective.

Optimizing the KL objective is faster than the KSD objective. This is because training the MLP classifier in Eq. (4) has cost linear in the dimension of the input, while calculating the Stein kernel in Eq. (13) has quadratic cost. However, training the MLP classifier may still be non-trivial in some instances, as we show in Appendix E. Furthermore, at a large scale the performance of the classifier may suffer in quality, if kept efficient, due to the intractability of computing the KL divergence [109]. In contrast, the kernelized method is sample efficient at all scales. However, for the small problem sizes we use here, the adversarial method gives better results on average, so we choose it as the primary method in all subsequent experiments.

B. Continuous observed variables and amortization with a hidden Markov model

In our second experiment, we simulate the adversarial VI method on the HMM in Fig. 1(b). The purpose was to

demonstrate two interesting features: continuous observed variables and amortization. We set up the HMM for $T = 8$ time steps [white circles in Fig. 1(b)], each represented by a Bernoulli latent variable with conditional dependency

$$z_1 \sim \mathcal{B}(\tfrac{1}{2}),$$

$$z_t \sim \begin{cases} \mathcal{B}(\tfrac{1}{3}) & \text{if } z_{t-1} = 0, \\ \mathcal{B}(\tfrac{2}{3}) & \text{if } z_{t-1} = 1. \end{cases} \quad (15)$$

These represent the unknown “regime” at time t . The regime affects how the observable data are generated. We use Gaussian observed variables [green circles in Fig. 1(b)] whose mean and standard deviation depend on the latent variables as

$$x_t \sim \begin{cases} \mathcal{N}(0, 1) & \text{if } z_t = 0, \\ \mathcal{N}(1, \tfrac{1}{2}) & \text{if } z_t = 1. \end{cases} \quad (16)$$

We sample two time series observations $\mathbf{x}^{(1)}, \mathbf{x}^{(2)}$ from the HMM and take this to be our data distribution. These time series are shown by green dots in Fig. 7. Instead of fitting two approximate posteriors *separately*, we use a single Born machine with amortization $|\psi(\boldsymbol{\theta}, \mathbf{x})\rangle$. We used the ansatz in Fig. 5 for eight qubits with a state preparation layer $S(\mathbf{x}) = \bigotimes_{t=1}^8 R_x(x_t)$. In practice, this encoding step can be done under more careful considerations [110–112]. Parameters $\boldsymbol{\theta}$ are initialized to small values at random. We optimize the KL objective and a MLP classifier with 16 inputs, 24 rectified linear units, and one sigmoid unit, and the system is trained for 3000 epochs. The learning rates are set to be 0.006 for the Born machine and 0.03 for the MLP. The Born machine used 100 samples to estimate each expectation value, the MLP used 100 samples from each distribution and minibatches of size ten.

The histograms in Fig. 7 show the ten most probable configurations of latent variables for the true posterior along with the probabilities assigned by the Born machine. Conditioning on data point $\mathbf{x}^{(1)}$, the inferred most likely explanation is $|01100011\rangle$ (i.e., the mode of the Born machine). This coincides with the true posterior mode. For $\mathbf{x}^{(2)}$, the inferred mode is $|10001000\rangle$, which differs from the true posterior mode $|10000000\rangle$ by a single bit. The regime switching has therefore been modeled with high accuracy.

Rather than focusing on the mode, one can make use of the whole distribution to estimate some quantity of interest. This is done by taking samples from the Born machine and using them in a Monte Carlo estimate for such a quantity. For example, we could predict the expected value of the next latent variable z_{T+1} given the available observations. For data point $\mathbf{x}^{(1)}$, this entails the estimation of

$$\mathbb{E}_{z_T \sim q_{\boldsymbol{\theta}}(z_T | \mathbf{x}^{(1)})} \mathbb{E}_{z_{T+1} \sim p(z_{T+1} | z_T)} [z_{T+1}]. \quad (17)$$

FIG. 7. Truncated, ordered histograms of the posteriors for two observed samples (a) $\mathbf{x}^{(1)}$ and (b) $\mathbf{x}^{(2)}$ of the hidden Markov model in Eqs. (15)–(16). The histograms are sorted by probability of the true posterior. The blue bars are the probabilities of the corresponding approximate posterior. The x axis shows the latent state for each bar and the observed data point $\mathbf{x}$. The lower panels are the time series of the data (c) $\mathbf{x}^{(1)}$ and (d) $\mathbf{x}^{(2)}$, as well as the corresponding modes of the true posterior and Born machine posterior as indicated with stars in the upper panel. (a) Posterior histogram given data $\mathbf{x}^{(1)}$. (b) Posterior histogram given data $\mathbf{x}^{(2)}$. (c) Time series and posterior modes for $\mathbf{x}^{(1)}$. (d) Time series and posterior modes for $\mathbf{x}^{(2)}$.

To conclude, we mention that inference in this HMM can be performed exactly and efficiently with classical algorithms. The most likely explanation for the observed data can be found with the Viterbi algorithm in time $O(T|S|^2)$, where T is the length of the sequence and $|S|$ is the size of the unobserved variables (e.g., $|S| = 2$ for Bernoulli variables). This is not the case for more general models. For example, a factorial HMM has multiple independent chains of unobserved variables [113]. Exact inference costs $O(TM|S|^{M+1})$, where M is number of chains, and is typically replaced by approximate inference [113,114]. Our VI methods are generic and apply to factorial and other HMMs without changes.

C. Demonstration on IBMQ with the “lung cancer” Bayesian network

As a final experiment, we test the performance on real quantum hardware using the slightly more complex “lung cancer” network [53], specifically using the five-qubit `ibmq_rome` quantum processor. We access this device using `PyQuil` [115] and `tket` [116].

The lung cancer network (also called the “Asia” network) is an example of a medical diagnosis Bayesian

network, and is illustrated by Fig. 1(c). This network was chosen since it is small enough to fit without adaptation onto the quantum device, but also comes closer to a “real-world” example. A more complicated version with extra variables can be found in Ref. [117]. We note that the version we use is slightly modified from that of Ref. [53] for numerical reasons.

The network has eight variables: two “symptoms” for whether the patient presented with dyspnoea (D) (shortness of breath) or had a positive x ray (X); four possible “diseases” causing the symptoms bronchitis (B), tuberculosis (T), lung cancer (L), or an “illness” (I) (which could be either tuberculosis or lung cancer or something other than bronchitis); and two possible “risk factors” for whether the patient had traveled to Asia (A) or whether they had a history of smoking (S).

Based on the graph structure in Fig. 1(c), the distribution over the variables, $p(A, T, S, L, I, X, B, D)$ can be factorized as

$$p(A)p(T|A)p(X|I)p(I|T, L)p(D|B, I)p(B|S)p(L|S)p(S). \quad (18)$$

FIG. 8. Histograms of true versus the learned posteriors with a simulated and hardware trained Born machine for the “lung cancer” network in Fig. 1(c). Here we use the `ibmq_rome` quantum processor (whose connectivity is shown in the inset). We condition on $X = \text{fa}$, $D = \text{fa}$, $I = \text{tr}$. The x axis shows the configuration of observed (gray represents X, D, I) and unobserved variables (blue represents A, S, T, L, B) corresponding to each probability (filled circles = tr , empty circles = fa). For hardware results, the histogram is generated using 1024 samples.

In Appendix F we show the explicit probability table we use (Table I) for completeness. Modifying an illustrative example of a potential “real-world” use case in Ref. [53], a patient may present in a clinic with an “illness” ($I = \text{tr}$) but no shortness of breath ($D = \text{fa}$) as a symptom. Furthermore, an x ray has revealed a negative result ($X = \text{fa}$). However, we have no patient history and we do not know which underlying disease is actually present. As such, in the experiment, we condition on having observed the “evidence” variables, $\mathbf{x}$: X, D , and I . The remaining five are the latent variables, $\mathbf{z}$, so we will require five qubits.

In Fig. 8, we show the results when using the five-qubit `ibmq_rome` quantum processor, which has a linear topology. The topology is convenient since it introduces no major overheads when compiling from our ansatz in Fig. 5. We plot the true posterior versus the one learned by the Born machine both simulated, and on the quantum processor using the best parameters found (after about 400 epochs simulated and about 50 epochs on the processor). To train in both cases, we use the same classifier as in Sec. VB with five inputs and ten hidden rectified linear units using the KL objective, which we observed to give the best results. We employ a two-layer ansatz in Fig. 5 ($L = 2$), and 1024 shots are taken from the Born machine in all cases. We observe that the simulated Born machine is able to learn the true posterior very well with these parameters, but the performance suffers on the hardware. That being said, the trained hardware model is still able to pick three out of four highest probability configurations of the network (shown on the x axis of Fig. 8). The hardware results could likely be improved with error mitigation techniques.

VI. DISCUSSION

We present two variational inference methods that can exploit highly expressive approximate posteriors given by quantum models. The first method is based on minimizing

the Kullback-Leibler divergence to the true posterior and relies on a classifier that estimates probability ratios. The resulting adversarial training may be challenging due to the large number of hyperparameters and well-known stability issues [118,119]. On the other hand, it only requires the ability to (i) sample from the prior $p(\mathbf{z})$, (ii) sample from the Born machine $q_\theta(\mathbf{z}|\mathbf{x})$, and (iii) calculate the likelihood $p(\mathbf{x}|\mathbf{z})$. We can apply the method as is when the prior is implicit. If the likelihood is also implicit, we can reformulate Eq. (3) in *joint contrastive* form [100] and apply the method with minor changes. This opens the possibility to train a model where both the generative and inference processes are described by Born machines (i.e., replacing the classical model in Fig. 2 with a quantum one). This model can be thought of as an alternative way to implement the quantum-assisted Helmholtz machine [15] and the quantum variational autoencoder [16].

We also present a second VI method based on the kernelized Stein discrepancy. A limitation of this approach is that it requires explicit priors and likelihoods. On the other hand, it provides plenty of flexibility in the choice of kernel. We use a generic Hamming kernel to compute similarities between bit strings. For VI on graphical models, a graph-based kernel [120] that takes into account the structure could yield improved results. Moreover, one could attempt a demonstration of quantum advantage in VI using classically intractable kernels arising from quantum circuits [39–41].

Another interesting extension could be approximating the posterior with a Born machine with traced out additional qubits, i.e., locally purified states. This may allow trading a larger number of qubits for reduced circuit depth at constant expressivity [86].

We present a few examples with direct application to financial, medical, and other domains. Another area of potential application of our methods is natural language processing. Quantum circuits have been proposed as a

way to encode and compose words to form meanings [121–123]. Question answering with these models could be phrased as variational inference on quantum computers where our methods fit naturally.

ACKNOWLEDGMENTS

This research was funded by Cambridge Quantum Computing. We gratefully acknowledge the cloud computing resources received from the “Microsoft for Startups” program. We acknowledge the use of IBM Quantum Services for this work. We thank Bert Kappen, Konstantinos Meichanetzidis, Robin Lorenz, and David Amaro for helpful discussions.

APPENDIX A: QUANTUM ADVANTAGE IN INFERENCE

The intuitive reason for why approximate inference may be a suitable candidate for a quantum advantage is similar to that which has already been studied for the related problem of generative modeling using quantum computers [42–45,85]. Here, the primary argument for quantum advantage reduces to complexity theoretic arguments for the hardness of sampling from particular probability distributions by any efficient classical means. At one end of the spectrum, the result of Ref. [44] is based on the classical hardness of the discrete logarithm problem and is not suitable for near-term quantum computers. On the other end, the authors of Refs. [43,85] leveraged arguments from “quantum computational supremacy” to show how Born machines are more expressive than their classical counterparts. Specifically, assuming the noncollapse of the polynomial hierarchy, the Born machine is able to represent IQP time computations [81].

A similar argument can be made for approximate inference. The key distinction in this case over the previous arguments for generative modeling is that now we must construct a conditional distribution, the posterior $p(\mathbf{z}|\mathbf{x})$, which at least requires a Born machine to *represent* it. For an IQP circuit, probability amplitudes are given by partition functions of complex Ising models [88,89]. Thus, for a demonstration of quantum advantage in inference, we may seek posterior distributions of the form $p(\mathbf{z}|\mathbf{x}) \propto |\text{tr}(e^{-iH(\mathbf{z},\mathbf{x})})|^2$. Here the Ising Hamiltonian $H(\mathbf{z}, \mathbf{x})$ is diagonal and its coefficients depend on both $\mathbf{z}$ and $\mathbf{x}$. Alternative “quantum computational supremacy” posteriors could arise from boson sampling [82] and random circuits [83,84]. In the latter case, experimental validations have been performed [124,125], demonstrating the *physical* existence of candidate distributions. The former experiment [124] used 53 qubits in the *Sycamore* quantum chip with 20 “cycles,” where a cycle is a layer of random single-qubit gates interleaved with alternating two-qubit gates. Notably, a cycle is not conceptually very different from a layer in our hardware-efficient ansatz in

Fig. 5 in the main text. The average (isolated) single-(two-)qubit gate error is estimated to be 0.13% (0.36%) in this experiment, which is sufficient to generate a cross-entropy (a measure of distribution “quality,” related to the KL divergence) benchmarking fidelity of 0.1%. The latter [125] repeated this experiment by using 56 qubits from the *Zuchongzhi* processor, again with 20 cycles. In this case, the single- (two-)qubit gates carried an average error of 0.14% (0.59%) to achieve a fidelity of approximately 0.07%. Clearly, quantum processors exist that can express classically intractable distributions, and so are also likely capable of representing intractable posterior distributions for suitably constructed examples. In all cases, the posterior shall be constructed nontrivially from the prior $p(\mathbf{z})$ and likelihood $p(\mathbf{x}|\mathbf{z})$ via Bayes’ theorem. We leave this construction for future work. As a final comment, it is also important to note here that we are discussing only *representational power* and not *learnability*, which is a separate issue. It is still unknown whether a distribution family (*a posterior* or otherwise) exists that is learnable by (near-term) quantum means, but not classically. See Refs. [43,44] for discussions of this point.

APPENDIX B: GRADIENTS FOR THE ADVERSARIAL METHOD

In the adversarial VI method we optimize the objective functions given in Eqs. (4) and (7) via gradient ascent and descent, respectively. Assuming that the optimal discriminator $d_\phi = d^*$ has been found, the objective in Eq. (7) becomes

$$\mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \mathbb{E}_{\mathbf{z} \sim q_\theta(\mathbf{z}|\mathbf{x})} \{\text{logit}[d^*(\mathbf{x}, \mathbf{z})] - \log p(\mathbf{x}|\mathbf{z})\}. \quad (\text{B1})$$

While d^* implicitly depends on θ , the gradient of its expectation vanishes. To see this, recall that $\text{logit}[d^*(\mathbf{z}, \mathbf{x})] = \log q_\theta(\mathbf{z}|\mathbf{x}) - \log p(\mathbf{z})$. Fixing $\mathbf{x}$ and taking the partial derivative for the j th parameter,

$$\begin{aligned} \partial_{\theta_j} \sum_{\mathbf{z}} q_\theta(\mathbf{z}|\mathbf{x}) \log q_\theta(\mathbf{z}|\mathbf{x}) &= \sum_{\mathbf{z}} \partial_{\theta_j} [q_\theta(\mathbf{z}|\mathbf{x})] \log [q_\theta(\mathbf{z}|\mathbf{x})] \\ &\quad + \sum_{\mathbf{z}} q_\theta(\mathbf{z}|\mathbf{x}) \partial_{\theta_j} [\log q_\theta(\mathbf{z}|\mathbf{x})]. \end{aligned} \quad (\text{B2})$$

The second term is zero since

$$\begin{aligned} \sum_{\mathbf{z}} q_\theta(\mathbf{z}|\mathbf{x}) \partial_{\theta_j} [\log q_\theta(\mathbf{z}|\mathbf{x})] &= \sum_{\mathbf{z}} \partial_{\theta_j} q_\theta(\mathbf{z}|\mathbf{x}) \\ &= \partial_{\theta_j} \sum_{\mathbf{z}} q_\theta(\mathbf{z}|\mathbf{x}) \\ &= \partial_{\theta_j} 1 \\ &= 0. \end{aligned} \quad (\text{B3})$$

The approximate posterior in Eq. (B2) is given by the expectations of observables $O_{\mathbf{z}} = |\mathbf{z}\rangle\langle\mathbf{z}|$, i.e., $q_\theta(\mathbf{z}|\mathbf{x}) =$

$\langle O_z \rangle_{\theta, \mathbf{x}}$. For a circuit parameterized by gates $U(\theta_j) = \exp[-i(\theta_j/2)H_j]$, where the Hermitian generator H_j has eigenvalues ± 1 (e.g., single-qubit rotations), the derivative of the first term in Eq. (B2) can be evaluated with the parameter shift rule [126–128]:

$$\partial_{\theta_j} q_{\theta}(\mathbf{z}|\mathbf{x}) = \frac{q_{\theta_j^+}(\mathbf{z}|\mathbf{x}) - q_{\theta_j^-}(\mathbf{z}|\mathbf{x})}{2} \quad (\text{B4})$$

with $\theta_j^{\pm} = \theta \pm (\pi/2)\mathbf{e}_j$ and $\mathbf{e}_j$ a unit vector in the j th direction. In summary, the partial derivative of the objective with respect to θ_j given the optimal discriminator is

$$\begin{aligned} \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} & \left(\mathbb{E}_{\mathbf{z} \sim q_{\theta_j^+}(\mathbf{z}|\mathbf{x})} \{ \text{logit}[d^*(\mathbf{x}, \mathbf{z})] - \log p(\mathbf{x}|\mathbf{z}) \} \right. \\ & \left. - \mathbb{E}_{\mathbf{z} \sim q_{\theta_j^-}(\mathbf{z}|\mathbf{x})} \{ \text{logit}[d^*(\mathbf{x}, \mathbf{z})] - \log p(\mathbf{x}|\mathbf{z}) \} \right). \end{aligned} \quad (\text{B5})$$

APPENDIX C: THE STEIN OPERATOR

Here we prove that $(O^p f)(\mathbf{z}) = s_p(\mathbf{x}, \mathbf{z})^T f(\mathbf{z}) - \text{tr}[\Delta f(\mathbf{z})]$ is a Stein operator for any function $f: \{0, 1\}^n \rightarrow \mathbb{R}^n$ and probability mass function $p(\mathbf{z}|\mathbf{x}) > 0$ on $\{0, 1\}^n$. To do so, we need to show that its expectation under the true posterior vanishes. Taking the expectation with respect to the true posterior we have

$$\begin{aligned} \mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{x})} [(O^p f)(\mathbf{z})] &= \sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) \frac{\Delta p(\mathbf{x}, \mathbf{z})^T}{p(\mathbf{x}, \mathbf{z})} f(\mathbf{z}) \\ &- \sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) \text{tr}[\Delta f(\mathbf{z})]. \end{aligned} \quad (\text{C1})$$

The first summation is

$$\begin{aligned} \sum_{\mathbf{z} \in \{0, 1\}^n} \frac{1}{p(\mathbf{x})} \Delta p(\mathbf{x}, \mathbf{z})^T f(\mathbf{z}) \\ = \sum_{\mathbf{z} \in \{0, 1\}^n} \sum_{i=1}^n [p(\mathbf{z}|\mathbf{x}) - p(\neg_i \mathbf{z}|\mathbf{x})] f_i(\mathbf{z}), \end{aligned} \quad (\text{C2})$$

while the second summation is

$$\begin{aligned} \sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) \text{tr}[\Delta f(\mathbf{z})] \\ = \sum_{\mathbf{z} \in \{0, 1\}^n} \sum_{i=1}^n p(\mathbf{z}|\mathbf{x}) [f_i(\mathbf{z}) - f_i(\neg_i \mathbf{z})]. \end{aligned} \quad (\text{C3})$$

Subtracting (C3) from (C2) we have

$$\begin{aligned} \sum_{i=1}^n \left[\sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) f_i(\mathbf{z}) - \sum_{\mathbf{z} \in \{0, 1\}^n} p(\neg_i \mathbf{z}|\mathbf{x}) f_i(\mathbf{z}) \right. \\ \left. - \sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) f_i(\mathbf{z}) + \sum_{\mathbf{z} \in \{0, 1\}^n} p(\mathbf{z}|\mathbf{x}) f_i(\neg_i \mathbf{z}) \right] = 0. \end{aligned} \quad (\text{C4})$$

The first and third terms cancel. The second and fourth terms are identical and also cancel. This can be seen by substituting $\mathbf{y}_i = \neg_i \mathbf{z}$ into one of the summations and noting that $\mathbf{z} = \neg_i \mathbf{y}_i$. Thus, the operator O^p satisfies Stein's identity $\mathbb{E}_{\mathbf{z} \sim p(\mathbf{z}|\mathbf{x})} [(O^p f)(\mathbf{z})] = 0$ for any function f as defined above, and is a valid Stein operator.

APPENDIX D: GRADIENTS FOR THE KERNELIZED METHOD

We wish to optimize the kernelized Stein objective, Eq. (14), via gradient descent. This gradient was derived in Ref. [43] and is given by

$$\begin{aligned} \partial_{\theta_j} \mathcal{L}_{\text{KSD}}(\boldsymbol{\theta}) &= \partial_{\theta_j} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]} \\ &= \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \partial_{\theta_j} \sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]} \\ &= \frac{1}{2} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \left[\frac{1}{\sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]}} \{ \partial_{\theta_j} \mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})] \} \right] \\ &= \frac{1}{4} \mathbb{E}_{\mathbf{x} \sim p_{\mathcal{D}}(\mathbf{x})} \left[\frac{1}{\sqrt{\mathbb{E}_{\mathbf{z}, \mathbf{z}' \sim q_{\theta}(\mathbf{z}|\mathbf{x})} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})]}} \left\{ \mathbb{E}_{\mathbf{z} \sim q_{\theta}, \mathbf{z}' \sim q_{\theta_j^+}} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim q_{\theta}, \mathbf{z}' \sim q_{\theta_j^-}} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})] \right. \right. \\ &\quad \left. \left. + \mathbb{E}_{\mathbf{z} \sim q_{\theta_j^+}, \mathbf{z}' \sim q_{\theta}} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim q_{\theta_j^-}, \mathbf{z}' \sim q_{\theta}} [\kappa_p(\mathbf{z}, \mathbf{z}'|\mathbf{x})] \right\} \right]. \end{aligned} \quad (\text{D1})$$

The last equality [suppressing the dependence on the variables, $\mathbf{z}, \mathbf{x}$: $q_\theta := q_\theta(\mathbf{z}|\mathbf{x})$] in Eq. (D1) follows from the parameter shift rule, Eq. (B4), for the gradient of the output probabilities of the Born machine circuit, under the same assumptions as in Appendix B. The relatively simple form of the gradient is due to the fact that the Stein kernel, κ_p , does not depend on the variational distribution, $q_\theta(\mathbf{z}|\mathbf{x})$. This gradient can be estimated in a similar way to the KSD itself, by sampling from the original Born machine, $q_\theta(\mathbf{z}|\mathbf{x})$, along with its parameter shifted versions, $q_{\theta_j^\pm}(\mathbf{z}|\mathbf{x})$ for each parameter j .

APPENDIX E: LEARNING CURVES FOR THE ADVERSARIAL METHOD

In this section we provide examples of successful and unsuccessful VI using the KL objective and adversarial methods. We inspect the random instances used in the “sprinkler” network experiment in Sec. V A. We cherry-pick two out of the 30 instances where a one-layer Born machine is employed. Figure 9(a) shows a successful experiment. Here the MLP classifier closely tracks the ideal classifier, providing a good signal for the Born machine. The Born machine is able to minimize its loss, which in turn leads to approximately 0 total variation distance. Figure 9(b) shows an unsuccessful experiment. Here, the MLP tracks the ideal classifier, but the Born machine is not able to find a “direction” that minimizes the loss (see epochs 400 to 1000). We verified that a more

powerful two-layer Born machine performs much better for this instance (not shown).

In these figures gray lines correspond to quantities that can be estimated from samples, while blue and magenta lines represent exact quantities that are not practical to compute. It is challenging to assess the training looking at the gray lines. This discussion is to emphasize the need for new techniques to assess the models. These would provide an early stopping criterion for unsuccessful experiments as well as a validation criterion for successful ones. This is a common problem among adversarial approaches.

APPENDIX F: PROBABILITY TABLE FOR THE “LUNG CANCER” NETWORK

We derived our VI methods under the assumption of nonzero posterior probabilities. In tabulated Bayesian networks, however, zero posterior probabilities may occur. The original “lung cancer” network [53] contains a variable “either” (meaning that either lung cancer or tuberculosis are present). This variable is a deterministic OR function of its parent variables and yields zero probabilities in some posterior distributions. In order to avoid numerical issues arising from zero probabilities, we replace “either” by “illness,” which may be true even if both parent variables “lung cancer” and “tuberculosis” are false. In practice, this is done by adding a small $\epsilon > 0$ to the entries of the “either” table and then renormalizing. This simple technique is known as additive smoothing, or Laplace

FIG. 9. Examples of successful (a) and unsuccessful (b) VI on the “sprinkler” Bayesian network in Sec. V A using the KL objective with the adversarial method. (a) Successful example. (b) Unsuccessful example.

TABLE I. Probability table corresponding to the network in Fig. 1(c).

Variable	Probabilities	Variable	Probabilities
Asia (A)	$p(A = \text{tr}) = 0.01$ $p(A = \text{fa}) = 0.99$	Smoking (S)	$p(S = \text{tr}) = 0.5$ $p(S = \text{fa}) = 0.5$
Tuberculosis (T)	$p(T = \text{tr} A = \text{tr}) = 0.05$ $p(T = \text{tr} A = \text{fa}) = 0.01$ $p(T = \text{fa} A = \text{tr}) = 0.95$ $p(T = \text{fa} A = \text{fa}) = 0.99$	Illness (I)	$p(I = \text{tr} L = \text{tr}, T = \text{tr}) = 0.95$ $p(I = \text{tr} L = \text{tr}, T = \text{fa}) = 0.95$ $p(I = \text{tr} L = \text{fa}, T = \text{tr}) = 0.95$ $p(I = \text{tr} L = \text{fa}, T = \text{fa}) = 0.05$ $p(I = \text{fa} L = \text{tr}, T = \text{tr}) = 0.05$ $p(I = \text{fa} L = \text{tr}, T = \text{fa}) = 0.05$ $p(I = \text{fa} L = \text{fa}, T = \text{tr}) = 0.05$ $p(I = \text{fa} L = \text{fa}, T = \text{fa}) = 0.95$
Lung cancer (L)	$p(L = \text{tr} S = \text{tr}) = 0.1$ $p(L = \text{tr} S = \text{fa}) = 0.01$ $p(L = \text{fa} S = \text{tr}) = 0.9$ $p(L = \text{fa} S = \text{fa}) = 0.99$	Dyspnoea (D)	$p(D = \text{tr} B = \text{tr}, I = \text{tr}) = 0.9$ $p(D = \text{tr} B = \text{tr}, I = \text{fa}) = 0.8$ $p(D = \text{tr} B = \text{fa}, I = \text{tr}) = 0.7$ $p(D = \text{tr} B = \text{fa}, I = \text{fa}) = 0.1$ $p(D = \text{fa} B = \text{tr}, I = \text{tr}) = 0.1$ $p(D = \text{fa} B = \text{tr}, I = \text{fa}) = 0.2$ $p(D = \text{fa} B = \text{fa}, I = \text{tr}) = 0.3$ $p(D = \text{fa} B = \text{fa}, I = \text{fa}) = 0.9$
Bronchitis (B)	$p(B = \text{tr} S = \text{tr}) = 0.6$ $p(B = \text{tr} S = \text{fa}) = 0.3$ $p(B = \text{fa} S = \text{tr}) = 0.4$ $p(B = \text{fa} S = \text{fa}) = 0.7$		
X ray (X)	$p(X = \text{tr} I = \text{tr}) = 0.98$ $p(X = \text{tr} I = \text{fa}) = 0.05$ $p(X = \text{fa} I = \text{tr}) = 0.02$ $p(X = \text{fa} I = \text{fa}) = 0.95$		

smoothing. Table I shows the modified probability table for the network in Fig. 1(c).

An alternative approach to deal with a deterministic variable that implements the OR function is to marginalize its parent variables. It can be verified that the entries of the resulting probability table are nonzero (unless the parent variables are also deterministic, in which case one marginalizes these as well). One can perform variational inference on this modified network with less variables.

- [1] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, 2009).
- [2] Q. Morris, in *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence (UAI2001)*, edited by J. Breese and D. Koller (Morgan Kaufmann Publishers Inc., Seattle, Washington, USA, 2001), p. 370, [ArXiv:1301.2295](#).
- [3] J. G. Richens, C. M. Lee, and S. Johri, Improving the accuracy of medical diagnosis with causal machine learning, *Nat. Commun.* **11**, 3923 (2020).
- [4] M. Maathuis, M. Drton, S. Lauritzen, and M. Wainwright, *Handbook of Graphical Models*, Chapman & Hall/CRC Handbooks of Modern Statistical Methods (CRC Press, Boca Raton, 2018).
- [5] W. Wiegand, B. Kappen, and W. Burgers, in *Interactive Collaborative Information Systems*, edited by R. Babuška and F. C. A. Groen (Springer, Berlin, Heidelberg, 2010), p. 547.
- [6] A. Denev, *Probabilistic Graphical Models: A New Way of Thinking in Financial Modelling* (Risk Books, 2015).

- [7] B. Cai, L. Huang, and M. Xie, Bayesian networks in fault diagnosis, *IEEE Trans. Ind. Inf.* **13**, 2227 (2017).
- [8] R. M. Neal, Technical Report CRG-TR-93-1, Department of Computer Science, University of Toronto, Toronto, 1993.
- [9] S. Brooks, A. Gelman, G. L. Jones, and X.-L. Meng, eds., *Handbook of Markov Chain Monte Carlo Methods* (Chapman & Hall/CRC, New York, NY, USA, 2011).
- [10] D. M. Blei, A. Kucukelbir, and J. D. McAuliffe, Variational inference: A review for statisticians, *J. Am. Stat. Assoc.* **112**, 859 (2017).
- [11] S. H. Adachi and M. P. Henderson, Application of quantum annealing to training of deep neural networks, [ArXiv:1510.06356](#) [quant-ph] (2015).
- [12] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, Estimation of effective temperatures in quantum annealers for sampling applications: A case study with possible applications in deep learning, *Phys. Rev. A* **94**, 022308 (2016).
- [13] M. Benedetti, J. Realpe-Gómez, R. Biswas, and A. Perdomo-Ortiz, Quantum-Assisted Learning of Hardware-Embedded Probabilistic Graphical Models, *Phys. Rev. X* **7**, 041052 (2017).
- [14] D. Korenkevych, Y. Xue, Z. Bian, F. Chudak, W. G. Macready, J. Rolfe, and E. Andriyash, Benchmarking quantum hardware for training of fully visible Boltzmann machines, [ArXiv:1611.04528](#) [quant-ph] (2016).
- [15] M. Benedetti, J. Realpe-Gómez, and A. Perdomo-Ortiz, Quantum-assisted helmholtz machines: A quantum-classical deep learning framework for industrial datasets in near-term devices, *Quantum Sci. Technol.* **3**, 034007 (2018).

- [16] A. Khoshaman, W. Vinci, B. Denis, E. Andriyash, H. Sadeghi, and M. H. Amin, Quantum variational autoencoder, *Quantum Sci. Technol.* **4**, 014001 (2018).
- [17] M. Wilson, T. Vandal, T. Hogg, and E. G. Rieffel, Quantum-assisted associative adversarial network: Applying quantum annealing in deep learning, *Quantum Mach. Intell.* **3**, 1 (2021).
- [18] G. H. Low, T. J. Yoder, and I. L. Chuang, Quantum inference on Bayesian networks, *Phys. Rev. A* **89**, 062315 (2014).
- [19] A. Montanaro, Quantum speedup of monte carlo methods, *Proc. R. Soc. A: Math., Phys. Eng. Sci.* **471**, 20150301 (2015).
- [20] A. N. Chowdhury and R. D. Somma, Quantum algorithms for gibbs sampling and hitting-time estimation, *Quant. Inf. Comp.* **17**, 0041 (2017). [ArXiv:1603.02940](#) [quant-ph].
- [21] P. Wittek and C. Gogolin, Quantum enhanced inference in Markov logic networks, *Sci. Rep.* **7** (2017).
- [22] J. Wu and T. H. Hsieh, Variational Thermal Quantum Simulation via Thermofield Double States, *Phys. Rev. Lett.* **123**, 220502 (2019).
- [23] G. Verdon, J. Marks, S. Nanda, S. Leichenauer, and J. Hidary, Quantum Hamiltonian-Based Models and the Variational Quantum Thermalizer Algorithm, [ArXiv:1910.02071](#) [quant-ph] (2019).
- [24] A. N. Chowdhury, G. H. Low, and N. Wiebe, A Variational Quantum Algorithm for Preparing Quantum Gibbs States, [ArXiv:2002.00055](#) [quant-ph] (2020).
- [25] Y. Wang, G. Li, and X. Wang, Variational quantum Gibbs state preparation with a truncated Taylor series, [ArXiv:2005.08797](#) [quant-ph] (2020).
- [26] Y. Shingu, Y. Seki, S. Watabe, S. Endo, Y. Matsuzaki, S. Kawabata, T. Nikuni, and H. Hakoshima, Boltzmann machine learning with a variational quantum algorithm, [ArXiv:2007.00876](#) [quant-ph] (2020).
- [27] I. Sato, K. Kurihara, S. Tanaka, H. Nakagawa, and S. Miyashita, in *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence UAI'09* (AUAI Press, Arlington, Virginia, USA, 2009), p. 479, [ArXiv:1408.2037](#).
- [28] B. O'Gorman, R. Babbush, A. Perdomo-Ortiz, A. Aspuru-Guzik, and V. Smelyanskiy, Bayesian network structure learning using quantum annealing, *Eur. Phys. J. Spec. Top.* **224**, 163 (2015).
- [29] H. Miyahara and Y. Sughiyama, Quantum extension of variational bayes inference, *Phys. Rev. A* **98**, 022330 (2018).
- [30] C. Zhang, J. Butepage, H. Kjellstrom, and S. Mandt, Advances in variational inference, [ArXiv:1711.05597](#) [cs.LG] (2018).
- [31] A. Abbas, D. Sutter, C. Zoufal, A. Lucchi, A. Figalli, and S. Woerner, The power of quantum neural networks, *Nat. Comput. Sci.* **1**, 403 (2021).
- [32] H.-Y. Huang, R. Kueng, and J. Preskill, Information-Theoretic Bounds on Quantum Advantage in Machine Learning, *Phys. Rev. Lett.* **126**, 190505 (2021).
- [33] K. Poland, K. Beer, and T. J. Osborne, No free lunch for quantum machine learning, [ArXiv:2003.14103](#) [quant-ph] (2020).
- [34] K. Sharma, M. Cerezo, Z. Holmes, L. Cincio, A. Sornborger, and P. J. Coles, Reformulation of the no-free-lunch theorem for entangled data sets, [ArXiv:2007.04900](#) [quant-ph] (2020).
- [35] R. A. Servedio and S. J. Gortler, Equivalences and separations between quantum and classical learnability, *SIAM J. Comput.* **33**, 1067 (2004).
- [36] S. Arunachalam and R. de Wolf, Guest column: A survey of quantum learning theory, *ACM SIGACT News* **48**, 41 (2017).
- [37] S. Arunachalam and R. d. Wolf, Optimal quantum sample complexity of learning algorithms, *J. Mach. Learn. Res.* **19**, 1 (2018).
- [38] Y. Liu, S. Arunachalam, and K. Temme, A rigorous and robust quantum speed-up in supervised machine learning, *Nat. Phys.* **17**, 1013 (2021).
- [39] 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* **567**, 209 (2019).
- [40] M. Schuld and N. Killoran, Quantum Machine Learning in Feature Hilbert Spaces, *Phys. Rev. Lett.* **122**, 040504 (2019).
- [41] H.-Y. Huang, M. Broughton, M. Mohseni, R. Babbush, S. Boixo, H. Neven, and J. R. McClean, Power of data in quantum machine learning, *Nat. Commun.* **12**, 2631 (2021).
- [42] X. Gao, Z.-Y. Zhang, and L.-M. Duan, A quantum machine learning algorithm based on generative models, *Sci. Adv.* **4**, eaat9004 (2018).
- [43] 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** (2020).
- [44] R. Sweke, J.-P. Seifert, D. Hangleiter, and J. Eisert, On the quantum versus classical learnability of discrete distributions, *Quantum* **5**, 417 (2021).
- [45] X. Gao, E. R. Anschuetz, S.-T. Wang, J. I. Cirac, and M. D. Lukin, Enhancing Generative Models via Quantum Correlations, [ArXiv:2101.08354](#) [quant-ph] (2021).
- [46] D. Ristè, M. P. da Silva, C. A. Ryan, A. W. Cross, A. D. Córcoles, J. A. Smolin, J. M. Gambetta, J. M. Chow, and B. R. Johnson, Demonstration of quantum advantage in machine learning, *npj Quantum Inf.* **3**, 1 (2017).
- [47] B. Coyle, M. Henderson, J. C. J. Le, N. Kumar, M. Paini, and E. Kashefi, Quantum versus classical generative modelling in finance, *Quantum Sci. Technol.* **6**, 024013 (2021).
- [48] J. Alcazar, V. Leyton-Ortega, and A. Perdomo-Ortiz, Classical versus quantum models in machine learning: Insights from a finance application, *Mach. Learning: Sci. Technol.* **1**, 035003 (2020).
- [49] S. Johri, S. Debnath, A. Mocherla, A. Singh, A. Prakash, J. Kim, and I. Kerenidis, Nearest centroid classification on a trapped ion quantum computer, [ArXiv:2012.04145](#) [quant-ph] (2020).
- [50] C. Ciliberto, M. Herbster, A. D. Ialongo, M. Pontil, A. Rocchetto, S. Severini, and L. Wossnig, Quantum machine learning: A classical perspective, *Proc. R. Soc. A: Math., Phys. Eng. Sci.* **474**, 20170551 (2018).
- [51] M. Benedetti, E. Lloyd, S. Sack, and M. Fiorentini, Parameterized quantum circuits as machine learning models, *Quantum Sci. Technol.* **4**, 043001 (2019).

- [52] L. Lamata, Quantum machine learning and quantum biomimetics: A perspective, *Mach. Learning: Sci. Technol.* **1**, 033002 (2020).
- [53] S. L. Lauritzen and D. J. Spiegelhalter, Local computations with probabilities on graphical structures and their application to expert systems, *J. R. Stat. Soc. Ser. B (Methodol.)* **50**, 157 (1988).
- [54] M. Kritzman, S. Page, and D. Turkington, Regime shifts: Implications for dynamic strategies (corrected), *Financ. Anal. J.* **68**, 22 (2012).
- [55] Z. Obermeyer, B. Powers, C. Vogeli, and S. Mullainathan, Dissecting racial bias in an algorithm used to manage the health of populations, *Science* **366**, 447 (2019).
- [56] J. Pearl, *Causality: Models, Reasoning and Inference* (Cambridge University Press, New York, NY, USA, 2009), 2nd ed.
- [57] R. McEliece, D. MacKay, and Jung-Fu Cheng, Turbo decoding as an instance of pearl's "belief propagation" algorithm, *IEEE J. Sel Areas Commun.* **16**, 140 (1998).
- [58] Dan Roth, On the hardness of approximate reasoning, *Artif. Intell.* **82**, 273 (1996).
- [59] G. F. Cooper, The computational complexity of probabilistic inference using bayesian belief networks, *Artif. Intell.* **42**, 393 (1990).
- [60] P. Dagum and M. Luby, Approximating probabilistic inference in Bayesian belief networks is NP-hard, *Artif. Intell.* **60**, 141 (1993).
- [61] R. Ranganath, S. Gerrish, and D. M. Blei, Black box variational inference, *ArXiv:1401.0118* [stat.ML] (2013).
- [62] D. P. Kingma and M. Welling, Auto-encoding variational Bayes, *ArXiv:1312.6114* [stat.ML] (2014).
- [63] D. Wingate and T. Weber, Automated variational inference in probabilistic programming, *ArXiv:1301.1299* [stat.ML] (2013).
- [64] S. Gershman, M. Hoffman, and D. Blei, Nonparametric variational inference, *ArXiv:1206.4665* [cs.LG] (2012).
- [65] D. J. Rezende and S. Mohamed, in *Proceedings of the 32nd International Conference on Machine Learning* (PMLR, 2015), Vol. 37, p. 1530, *ArXiv:1505.05770* [stat.ML].
- [66] R. J. Williams, Simple statistical gradient-following algorithms for connectionist reinforcement learning, *Mach. Learn.* **8**, 229 (1992).
- [67] C. J. Maddison, A. Mnih, and Y. W. Teh, The concrete distribution: A continuous relaxation of discrete random variables, *ArXiv:1611.00712* [cs.LG] (2017).
- [68] E. Jang, S. Gu, and B. Poole, Categorical reparameterization with gumbel-softmax, *ArXiv:1611.01144* [stat.ML] (2017).
- [69] S. Cheng, J. Chen, and L. Wang, Information perspective to probabilistic modeling: Boltzmann machines versus born machines, *Entropy* **20**, 583 (2018).
- [70] 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** (2019).
- [71] K. E. Hamilton, E. F. Dumitrescu, and R. C. Pooser, Generative model benchmarks for superconducting qubits, *Phys. Rev. A* **99**, 062323 (2019).
- [72] V. Leyton-Ortega, A. Perdomo-Ortiz, and O. Perdomo, Robust implementation of generative modeling with parametrized quantum circuits, *Quantum Mach. Intell.* **3**, 17 (2021).
- [73] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, Training of quantum circuits on a hybrid quantum computer, *Sci. Adv.* **5** (2019).
- [74] 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).
- [75] M. S. Rudolph, N. B. Toussaint, A. Katabarwa, S. Johri, B. Peropadre, and A. Perdomo-Ortiz, Generation of high-resolution handwritten digits with an ion-trap quantum computer, *ArXiv:2012.03924* [quant-ph] (2020).
- [76] I. Čepaitė, B. Coyle, and E. Kashefi, A continuous variable Born machine, *ArXiv:2011.00904* [quant-ph] (2020).
- [77] C. Zoufal, A. Lucchi, and S. Woerner, Quantum generative adversarial networks for learning and loading random distributions, *npj Quantum Inf.* **5** (2019).
- [78] D. Herr, B. Obert, and M. Rosenkranz, Anomaly detection with variational quantum generative adversarial networks, *Quantum Sci. Technol.* **6**, 045004 (2021).
- [79] Some literature (e.g., Ref. [77]) refers to Born machines as quantum generative adversarial networks (QGANs) when using an adversarial objective function. It is more appropriate to reserve the term QGAN to adversarial methods for quantum data (e.g., Refs. [129–131]).
- [80] A quantum computer cannot efficiently implement arbitrary circuits; thus, some Born machines are intractable even for a quantum computer.
- [81] M. J. Bremner, A. Montanaro, and D. J. Shepherd, Average-Case Complexity versus Approximate Simulation of Commuting Quantum Computations, *Phys. Rev. Lett.* **117**, 080501 (2016).
- [82] S. Aaronson and A. Arkhipov, in *Proceedings of the Forty-Third Annual ACM Symposium on Theory of Computing*, STOC'11 (Association for Computing Machinery, New York, NY, USA, 2011), p. 333.
- [83] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, Characterizing quantum supremacy in near-term devices, *Nat. Phys.* **14**, 595 (2018).
- [84] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, On the complexity and verification of quantum random circuit sampling, *Nat. Phys.* **15**, 159 (2019).
- [85] Y. Du, M.-H. Hsieh, T. Liu, and D. Tao, Expressive power of parametrized quantum circuits, *Phys. Rev. Res.* **2** (2020).
- [86] I. Glasser, R. Sweke, N. Pancotti, J. Eisert, and I. Cirac, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32, p. 1496, *ArXiv:1907.03741* [cs.LG].
- [87] S. Mohamed and B. Lakshminarayanan, Learning in implicit generative models, *ArXiv:1610.03483* [stat.ML] (2017).
- [88] G. D. las Cuevas, W. Dür, M. V. den Nest, and M. A. Martin-Delgado, Quantum algorithms for classical lattice models, *New J. Phys.* **13**, 093021 (2011).

- [89] K. Fujii and T. Morimae, Commuting quantum circuits and complexity of ising partition functions, *New J. Phys.* **19**, 033003 (2017).
- [90] Bias and variance of the estimates need to be analyzed and controlled, which is a challenge in itself.
- [91] J.-G. Liu and L. Wang, Differentiable learning of quantum circuit born machines, *Phys. Rev. A* **98**, 062324 (2018).
- [92] H. Situ, Z. He, Y. Wang, L. Li, and S. Zheng, Quantum generative adversarial network for generating discrete distribution, *Inf. Sci. (Ny)* **538**, 193 (2020).
- [93] J. Zeng, Y. Wu, J.-G. Liu, L. Wang, and J. Hu, Learning and inference on generative adversarial quantum circuits, *Phys. Rev. A* **99**, 052306 (2019).
- [94] S. E. Borujeni, S. Nannapaneni, N. H. Nguyen, E. C. Behrman, and J. E. Steck, Quantum circuit representation of Bayesian networks, *ArXiv:2004.14803* [quant-ph] (2020).
- [95] R. R. Tucci, Quantum Bayesian nets, *Int. J. Mod. Phys. B* **09**, 295 (1995).
- [96] R. Ranganath, J. Alotaar, D. Tran, and D. M. Blei, Operator variational inference, *ArXiv:1610.09033* [stat.ML] (2018).
- [97] B. K. Sriperumbudur, K. Fukumizu, A. Gretton, B. Schölkopf, and G. R. G. Lanckriet, On integral probability metrics, ϕ -divergences and binary classification, *ArXiv:0901.2698* [cs.IT] (2009).
- [98] D. J. Rezende, S. Mohamed, and D. Wierstra, Stochastic backpropagation and approximate inference in deep generative models, *ArXiv:1401.4082* [stat.ML] (2014).
- [99] L. Mescheder, S. Nowozin, and A. Geiger, Adversarial variational Bayes: Unifying variational autoencoders and generative adversarial networks, *ArXiv:1701.04722* [cs.LG] (2018).
- [100] F. Huszár, Variational inference using implicit distributions, *ArXiv:1702.08235* [stat.ML] (2017).
- [101] J. Yang, Q. Liu, V. Rao, and J. Neville, in *Proceedings of the 35th International Conference on Machine Learning*, Proceedings of Machine Learning Research, edited by J. Dy and A. Krause (PMLR, 2018), Vol. 80, p. 5561.
- [102] In contrast to Ref. [101] but in line with earlier work such as Ref. [132] we include the absolute value in the definition of the Stein discrepancy.
- [103] 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).
- [104] M. Cerezo, A. Sone, T. Volkoff, L. Cincio, and P. J. Coles, Cost function dependent barren plateaus in shallow parametrized quantum circuits, *Nat. Commun.* **12** (2021).
- [105] S. Wang, E. Fontana, M. Cerezo, K. Sharma, A. Sone, L. Cincio, and P. J. Coles, Noise-induced barren plateaus in variational quantum algorithms, *ArXiv:2007.14384* [quant-ph] (2021).
- [106] Z. Holmes, K. Sharma, M. Cerezo, and P. J. Coles, Connecting ansatz expressibility to gradient magnitudes and barren plateaus, *ArXiv:2101.02138* [quant-ph] (2021).
- [107] A. Pesah, M. Cerezo, S. Wang, T. Volkoff, A. T. Sornborger, and P. J. Coles, Absence of barren plateaus in quantum convolutional neural networks, *ArXiv:2011.02966* [quant-ph] (2020).
- [108] C. Zhao and X.-S. Gao, Analyzing the barren plateau phenomenon in training quantum neural network with the ZX-calculus, *Quantum* **5**, 466 (2021).
- [109] V. Kontonis, S. Liu, and C. Tzamos, Convergence and sample complexity of SGD in GANs, *ArXiv:2012.00732* [cs.LG] (2020).
- [110] R. LaRose and B. Coyle, Robust data encodings for quantum classifiers, *Phys. Rev. A* **102**, 032420 (2020).
- [111] M. Schuld, R. Sweke, and J. J. Meyer, Effect of data encoding on the expressive power of variational quantum-machine-learning models, *Phys. Rev. A* **103**, 032430 (2021).
- [112] A. Pérez-Salinas, A. Cervera-Lierta, E. Gil-Fuster, and J. I. Latorre, Data re-uploading for a universal quantum classifier, *Quantum* **4**, 226 (2020).
- [113] Z. Ghahramani and M. I. Jordan, Factorial hidden Markov models, *Mach. Learn.* **29**, 245 (1997).
- [114] J. Z. Kolter and T. Jaakkola, in *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, Proceedings of Machine Learning Research, edited by N. D. Lawrence and M. Girolami (PMLR, La Palma, Canary Islands, 2012), Vol. 22, p. 1472.
- [115] R. S. Smith, M. J. Curtis, and W. J. Zeng, A practical quantum instruction set architecture, *ArXiv:1608.03355* [quant-ph] (2016).
- [116] S. Sivarajah, S. Dilkes, A. Cowtan, W. Simmons, A. Edgington, and R. Duncan, tket: A retargetable compiler for NISQ devices, *Quantum Sci. Technol.* **6**, 014003 (2020).
- [117] D. Barber, *Bayesian Reasoning and Machine Learning* (Cambridge University Press, Cambridge, UK, 2012).
- [118] M. Arjovsky and L. Bottou, Towards principled methods for training generative adversarial networks, *ArXiv:1701.04862* [stat.ML] (2017).
- [119] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet, in *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, NIPS'18 (Curran Associates Inc., 2018), p. 698, *ArXiv:1711.10337* [stat.ML].
- [120] R. I. Kondor and J. D. Lafferty, in *Proceedings of the Nineteenth International Conference on Machine Learning*, ICML'02 (Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002), p. 315.
- [121] B. Coecke, G. de Felice, K. Meichanetzidis, and A. Toumi, Foundations for near-term quantum natural language processing, *ArXiv:2012.03755* [quant-ph] (2020).
- [122] K. Meichanetzidis, A. Toumi, G. de Felice, and B. Coecke, Grammar-aware question-answering on quantum computers, *ArXiv:2012.03756* [quant-ph] (2020).
- [123] R. Lorenz, A. Pearson, K. Meichanetzidis, D. Kartsaklis, and B. Coecke, Qnlp in practice: Running compositional models of meaning on a quantum computer, *ArXiv:2102.12846* [cs.CL] (2021).
- [124] F. Arute *et al.*, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505 (2019).
- [125] Y. Wu *et al.*, Strong quantum computational advantage using a superconducting quantum processor, *ArXiv:2106.14734* [quant-ph] (2021).

- [126] K. Mitarai, M. Negoro, M. Kitagawa, and K. Fujii, Quantum circuit learning, [Phys. Rev. A **98**, 032309 \(2018\)](#).
- [127] M. Schuld, V. Bergholm, C. Gogolin, J. Izaac, and N. Killoran, Evaluating analytic gradients on quantum hardware, [Phys. Rev. A **99**, 032331 \(2019\)](#).
- [128] L. Banchi and G. E. Crooks, Measuring analytic gradients of general quantum evolution with the stochastic parameter shift rule, [Quantum **5**, 386 \(2021\)](#).
- [129] P.-L. Dallaire-Demers and N. Killoran, Quantum generative adversarial networks, [Phys. Rev. A **98**, 012324 \(2018\)](#).
- [130] M. Benedetti, E. Grant, L. Wossnig, and S. Severini, Adversarial quantum circuit learning for pure state approximation, [New J. Phys. **21**, 043023 \(2019\)](#).
- [131] S. Chakrabarti, H. Yiming, T. Li, S. Feizi, and X. Wu, in *Advances in Neural Information Processing Systems*, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (Curran Associates, Inc., 2019), Vol. 32, p. 6781, [ArXiv:1911.00111](#).
- [132] J. Gorham and L. Mackey, Measuring Sample Quality with Stein's Method, [ArXiv:1506.03039](#) [cs, math, stat] (2018).