Variational Boson Sampling

Shiv Shankar[1] and Don Towsley[1]

University of Massachusetts
{sshankar,towsley}@cs.umass.edu

Abstract. Boson Samplers are near-term quantum devices based on photonic quantum technology, which can outperform classical computing systems. This paper takes a hybrid circuit learning approach to utilize boson samplers as a generative model called Variational Boson Sampling (VBS). VBS introduces an optimizable parametric structure into the evolution operator for boson sampling and uses the complete model as a variational ansatz. To simulate working with real quantum devices, we use gradient free-optimization methods to optimize the resultant circuit. We experiment with this framework for problems in optimization and generative modeling.
Index Terms— hybrid quantum-classical approach, quantum circuit learning, boson sampling, parameterized quantum circuit

1 Introduction

Quantum computing is a computational paradigm that utilizes properties of quantum systems such as superposition and entanglement for computational tasks (Nielsen and Chuang, 2002). Research in quantum computing has led to the discovery of quantum algorithms that achieve polynomial-time speedups compared to classical methods on specific tasks. For example, Grover (1996) described a sub-linear time algorithm for search in an unordered database. Similarly, Shor's algorithm (Shor, 1999) is capable of integer factorization in polynomial-time. Quantum techniques have also been proposed for data-science tasks such as regression fitting (Yu et al., 2019) and association rules mining (Yu et al., 2016).

While the prospect of significant speedup over classical computing remains the most impactful aspect of quantum computing, the algorithms mentioned above remain largely out of the reach of current quantum devices. Hence researchers have explored options that use near-term devices and can achieve quantum supremacy (Arute et al., 2019; Lund et al., 2017). Boson sampling (BS) (Aaronson and Arkhipov, 2011) is a strong candidate for experimental demonstration of quantum algorithmic supremacy.

However, while boson sampling (BS) provides an exponential quantum advantage and has been used for applications such as combinatorial optimization (Arrazola and Bromley, 2018), there is some debate on its demonstrable advantage in 'realistic' applications (Bromley et al., 2021; Oh et al., 2021). Our goal in this work is to provide examples of how BS can be used in ML applications, specifically as a generative model.

One effective approach for capturing quantum advantage is the so-called hybrid quantum-classical (HQC) approach (McClean et al., 2016). The HQC approach uses a combination of both quantum and classical resources. HQC based models have been getting recent traction and have been used for applications like supervised regression (Schuld et al., 2020; Yu et al., 2016), clustering (Otterbach et al., 2017) and combinatorial optimization (Moll et al., 2018).

This work utilizes a similar hybrid circuit approach to use boson samplers for generative models. Since the sampling distribution of a BS has an intractable dependence on the interferometer design, we propose a specific parameterization that allows exploration of the space of possible interferometers in a way that is both universal and amenable to near term devices. We then train networks following this design for a specific task by stochastic gradient-based minimization of the task loss. We conduct experiments with Ising models and image generation to evaluate the proposed scheme. The results show that this is a promising scheme to deploy boson sampling in practical tasks.

2 Preliminaries and Related Work

2.1 Hybrid Quantum-classical Approach

A hybrid quantum-classical algorithm consists of a classical computer and a quantum device running in a closed loop. Generally, the classical computer is used to preprocess the input data. Next, the data is transferred onto an initial quantum state that the quantum circuit acts upon. Finally, the output of the quantum device is evaluated against the desired output (usually on the classical computer). The resultant discrepancy between desired and actual output is then fed back into the system to adjust the quantum device to optimize the performance.

A standard HQC approach is based upon using a quantum circuit with parameterized gates. Such models, known as Parameterized Quantum Circuits (PQC), have demonstrated success in applications like supervised learning (Yu et al., 2016; Liu et al., 2019), generative modeling (Zeng et al., 2019) and algorithm learning (Morales et al., 2018) The PQC model most closely associated with this paper is the quantum circuit born machine (QCBM) (Liu and Wang, 2018a). Born machines can be used to efficiently generate fast mixing MCMC samples via projective measurement (Liu and Wang, 2018a). They are also capable of producing distributions that are classically intractable (Bouland et al., 2018).

2.2 Variational Quantum Born Machine

A QCBM uses a variational quantum circuit to encode and sample from the probability distribution of a classical dataset. Given a set of D independent and identically distributed samples $\{x_1, ..., x_D\}$ from a target probability distribution, the QCBM can be optimized to generate samples that approximate the unknown target probability distribution. As shown in Figure 8a, the QCBM takes the product state $|0\rangle^{\otimes n}$ and evolves it to the output state $|\psi_\phi\rangle$ by unitary transformation

$U(\phi)$. A common and universal circuit layout includes a series of single-qubit rotation gates and two-qubit CNOT gates. The rotation gates are parameterized with rotation parameters $\Phi = \{\phi_\alpha^d\}$ where d is the layer index and α indexes the gate in a layer.

For each rotation gate R^α at level d, the vector ϕ_α^d describe the three rotation parameters which describe the complete rotation gate by

$$R^\alpha(\phi_d) = R_z(\phi_{\alpha,x}^d)R_x(\phi_{\alpha,y}^d)R_z(\phi_{\alpha,z}^d)(1)$$

with $R_m(\phi) = exp(-i\frac{\phi\sigma_m}{2})$ and σ_m are the Pauli matrices. Samples can be obtained by measuring the output state $|\psi_\phi\rangle$ in the computational basis which will produce them by the probability $|\langle x|\psi_\phi\rangle|^2$

For machine learning applications such as classification, the output model distribution $p_\phi(x)$ is then optimized to minimize the negative log-likelihood of the observed data. The parameters are usually tuned via gradient descent though other methods can also be deployed Wang et al. (2019).

2.3 Boson Sampling

The BosonSampling (BS) problem refers to sampling outcomes from a linear optical network. The seminal paper of Aaronson and Arkhipov (2011) demonstrates how the simulation of the probability distribution of indistinguishable photons evolving in such a circuit is classically intractable. More specifically Aaronson and Arkhipov (2011) define a model where N isolated photons are sent through a m ($m > 2N$) mode linear-optical circuit/interferometer. The interferometer is described by a matrix $U \in \mathbb{U}(m)$, which transforms m input modes into m output modes. An example circuit is depicted in Figure 1b. Let $\bar{n} = |n_1, n_2, \dots n_m\rangle$ denote the output pattern with n_j photons in output j. The quantum state of the output photons is given by :

$$|\psi\rangle = \gamma_{\bar{n}} |n_1, n_2, \dots n_m\rangle$$
$$\gamma_{\hat{\bar{n}}} = \frac{\text{Perm}(U_S)}{\sqrt{\hat{\bar{n}}!}}$$

where $\hat{n}! = n_1!n_2! \dots n_m!$. U_S is the submatrix of U obtained by selecting the columns corresponding to input photons and rows corresponding to output photons. The probability of a certain output is given by

$$Pr(\bar{n}) = \frac{|\text{Perm}(U_S)|^2}{\bar{n}!}$$

Perm here refers to the permanent of the matrix. Computation of the permanent is #P-complete (Valiant, 1979), which makes exact sampling from such a circuit intractable. Further work (Morimae et al., 1998; Gogolin et al., 2013; Bremner et al., 2011) provided further connections between boson sampling and the polynomial hierarchy. Since a deterministic source of single photons as

described in Aaronson and Arkhipov (2011) is physically challenging; variants such as Lattice walk sampling (LWS) (Muraleedharan et al., 2019) and Gaussian boson sampling (GBS) (Lund et al., 2014; Hamilton et al., 2017) have been proposed. Recently, a quantum computing machine that uses Gaussian boson sampling was used to demonstrate quantum supremacy and was faster than the state-of-the-art classical supercomputers by a factor of 10^{14} (Zhong et al., 2020).

Recent research has also explored other applications of BS to solve diverse problems. Guerreschi (2015) demonstrated that molecular vibronic spectra could be efficiently generated using boson sampling. Arrazola and Bromley (2018) showed that boson sampling could be utilized for approximating the densest k-subgraph problem. Recently Huang et al. (2019) developed a quantum symmetric encryption scheme built on boson sampling. A quantum signature protocol using BS-based unitary operation (Gao et al., 2018) has been demonstrated by Feng et al. (2020). Bosonic techniques have also been successfully used for graph similarity-based tasks in machine learning (Schuld et al., 2020; Shankar and Towsley, 2020).

The work most related to our article is by Banchi et al. (2020). They show how under certain parameterization, the unbiased estimates of the gradients of the parameters for a GBS device can be obtained directly via measurements on the same device. Our work while also uses parameterized boson sampling is similar in spirit to the variational learning of a quantum Born machine (Liu and Wang, 2018b). Three other major differences from Banchi et al. (2020) are instead of Gaussian Boson sampling we focus on the classical boson sampling method, our approach is not restricted to the strict parameterizations and we can use our approach for more generic applications rather than simple graph based problems.

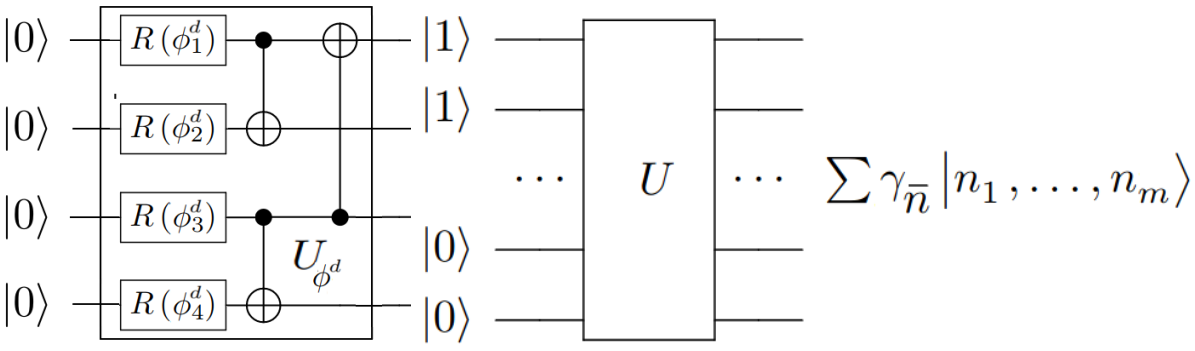

Fig. 1: Born Machine and Boson Sampling scheme. a) depicts a Variational QCBM circuit from Liu and Wang (2018b) with parameters ϕ^d b) depicts Boson Sampling Scheme with evolution matrix U

3 Variational Boson Sampling Circuits

This section presents a variational boson sampling (VBS) scheme that can be used as a generative model for fitting a target distribution. We describe the general VBS scheme before providing a flexible circuit design that remains efficient and scalable by limiting the parameter growth as a polynomial function of the circuit input size.

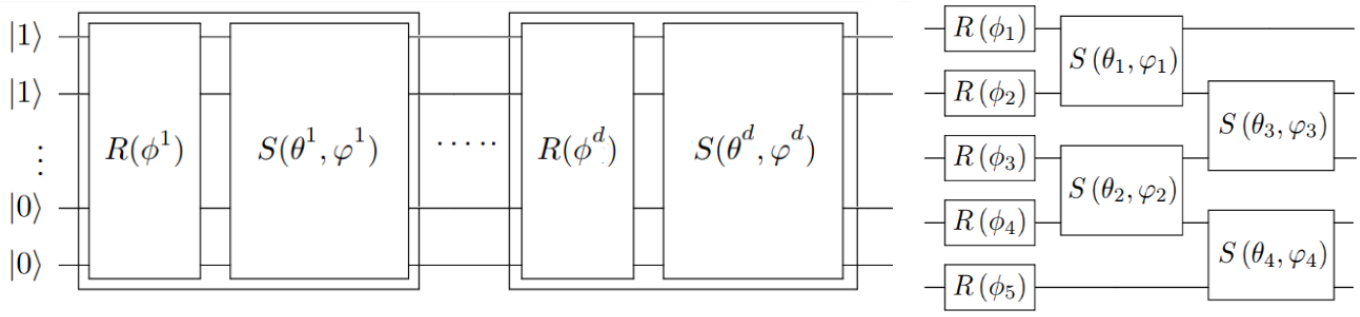

Fig. 2: Variational Boson Sampling scheme. (Left) The full circuit consists of layers $R(\phi^d)$ and $S(\theta^d, \varphi^d)$ blocks. (Right) A fixed VBS block analogous to the variational QCBM circuit of Figure 8a

Figure 2 depicts the schematic of the VBS model. The VBS scheme treats the quantum unitary in the BS scheme as a parametric function. This can be achieved by decomposing the original unitary matrix into parametric gates. Any unitary transform on n qubits can be decomposed into the product of multiple gates, each of which acts on at most two qubits (Nielsen and Chuang, 2002). Similarly, the $m \times m$ unitary transformation U in the BS model can also be decomposed into a combination of phase-shifters and beam-splitters (which act as single-qubit and two-qubit gates, respectively). Similar to a single qubit rotation gate, a phase shifter $R(\phi)$ acts only on a single-mode by multiplying the amplitude α_a of the corresponding mode by $e^{i\phi}$. A beam splitter $S(\theta, \varphi)$ on the other hand, acts on two modes as follows:

$$S(\theta, \varphi) = \begin{bmatrix} \cos\theta & -e^{-i\varphi}\sin(\theta) \\ e^{-i\varphi}\sin(\theta) & \cos\theta \end{bmatrix}$$

This decomposition has the added advantage that now, both the input and output modes of a VBS scheme can be coupled with a more standard PQC circuit for further computation. The final probability of the output $\bar{n}$ under this model is given by:

$$P_{\text{VBS}}(\bar{n}|\Omega) = \frac{|\text{Perm}(U_\Omega(\bar{n})|^2}{n_1! \cdots n_m!}. \tag{1}$$

Here Ω refers to all the parameters used to define the operator U_Ω. Specifically $\Omega = \{\boldsymbol{\phi}^k, \boldsymbol{\theta}^k, \boldsymbol{\varphi}^k\}$. The parameterized evolution operator U_Ω is given by $\prod_k (R(\boldsymbol{\phi}^k) S(\boldsymbol{\theta}^k, \boldsymbol{\varphi}^k))$, and $U_\Omega(\bar{n})$ corresponds to the submatrix of U_Ω obtained by selecting only the non-empty input columns and non-empty output rows from U_Ω.

Next, we present a fixed structure design that is both flexible and minimizes parameter explosion in the VBS scheme. We follow the design principle used by Liu and Wang (2018b), as this structure can be efficiently simulated via tensor networks. Each individual block of the VBS is a small-scale tensor network with a sparse, primarily diagonal structure. These networks can be composed together to form the larger unitary U_{VBS} of the entire circuit.

Similar to the design of Liu et al. (2019) (see Figure 8a), our proposed circuit is composed of blocks. Each block is composed of arrays of single-mode phase shifters and binary-mode beam splitters stacked together (see Figure 2 Left). First, the block applies an array of phase shifters $R(\boldsymbol{\phi_l})$ applied to all modes. The output is then presented to an entanglement layer composed of parametric beam splitters $S(\boldsymbol{\theta}_l, \boldsymbol{\varphi_l})$.

The individual $R(\phi)$ and $S(\theta, \varphi)$ blocks are composed of unary and binary gates. Each $R(\phi)$ block consists of an array of unary rotation gates/phase shifters that act on only one mode. Each $S(\theta, \varphi)$ layer consists of two arrays of beam splitter acting in a staggered way. An example of the combined composition of gates is presented in Figure 2 Right.

If the VBS circuit has L layers and m modes, then the total gate complexity of the circuit is mL phase shifters and mL beam splitters.

4 Experiments

Next, we try the aforementioned model for two tasks. The first task is to identify the ground states of an Ising Hamiltonian. Our experiments show that the VBS model can be trained to preferentially sample low-energy states from an Ising model. For the second task, we train a classically augmented VBS on the digits dataset (Alpaydin and Kaynak, 1998) to generate similar images. The results show that for comparable latent dimensions the VBS scheme is as expressive as Variational Autoencoders Kingma and Welling (2019).

4.1 Ising Model Optimization

An Ising model is essentially an energy model (or an unnormalized distribution) for which the score function (or the log-likelihood) is of the following form:

$$H(\bar{x}) = -\sum_i h_i x_i - \sum_{ij} J_{ij} x_i x_j, \tag{2}$$

where $\bar{x} = (x_1, x_2, \ldots, x_m)$ and $x_k = 0, 1$ i.e. $\bar{x}$ is a binary vector. We are interested in finding a model distribution that samples the state with the lowest energy

(also known as the ground state) with high probability. This is a challenging task, as finding such a state of a general $H(\bar{x})$ is NP-hard Lucas (2014). While similar to sampling from an energy model, this task is closer to optimization as we wish to find the lowest energy state. In this experiment the VBS parameters (Ω) are updated so that it samples the minimum energy configuration with high probability.

We follow the procedure of Banchi et al. (2020) and use our VBS model for predicting cliques. Since a boson sampling scheme produces an output with an integer number of bosons in different modes, the output of such a scheme can be thresholded to a vector $\bar{x}$ of binary variables for input to the Ising Hamiltonian. The training loss is given by:

$$\mathcal{L}(W) = \mathbb{E}_{\bar{x}\sim P_{\mathrm{VBS}}(.|\Omega)}\left[H(\bar{x})\right] \equiv \sum_{\bar{x}} H(\bar{x}) P_{\mathrm{VBS}}(\bar{x}|\Omega) \tag{3}$$

where $P_{\mathrm{VBS}}(\bar{x}|\Omega)$ is the distribution of Equation (1).

Similar to Banchi et al. (2020), we focus on the following Ising hamiltonian. Given a graph $G = (V, E)$ with vertex set V and edge set E and an integer K:

$$H_K(\bar{x}) = \lambda \left(K - \sum_{v\in V} x_v \right)^2 - \sum_{(u,v)\in E} x_u x_v \tag{4}$$

where λ is a positive number and x_v are binary variables. It is easy to prove (Lucas, 2014) that for $\lambda > K$ the above Hamiltonian has ground state energy $E = -\frac{K(K-1)}{2}$ if and only if there is a clique of size K in the graph G. To see this note that the second term $x_u x_v$ computes the number of edges between the set of nodes corresponding to the binary vector x and for a clique of size K will contribute $\frac{K(K-1)}{2}$. On the other hand the first term tries to keep the number of selected nodes to K. With a large enough λ, If the number of selected nodes becomes more than K, then the increase in the first term is enough to compensate for the reduction in the second term.

Following the procedure of Banchi et al. (2020) we use sampling to produce binary strings that corresponds to the ground state of the aforementioned Ising Hamiltonian. For each graph we set K as the size of the largest clique in the graph. Furthermore we set $\lambda = 2\Delta$ where Δ is the max degree of the graph. The success rate is estimated as the fraction of times that the correct bit pattern is sampled by the model in a 1000 samples, conditional on observing K output particles. Training is done using an estimation of the gradient using the REINFORCE algorithm (Williams, 1992), obtained with 200 samples per iteration.

We run VBS on the hamiltonian corresponding to the graphs experimented on by Banchi et al. (2020). These experiment are on the specific graphs depicted in 3 followed by a bunch of random graphs from the Erdos-Renyi and Barabasi Albert families (Figure 7). Figure 3(a) presents the training curve on a simple graph on 8 nodes with a clique size of 5. From the figure, it is clear that while the initial probability of sampling the ground state is low; it steadily increases as training progresses and is above 80% by the end. In Figure 3(b), a more challenging case

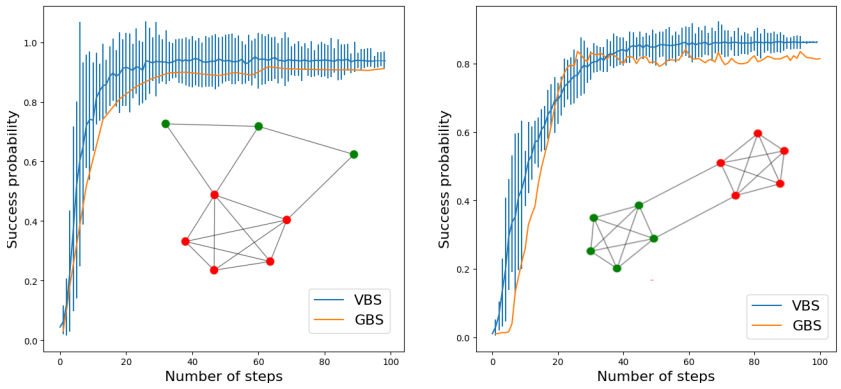

Fig. 3: Success rate of sampling the ground state of the Ising model over time for VBS and GBS. A max clique of size $K = 5$ is shown in red. In (a) there is a single max-clique, while in (b) there are two max-cliques. Training is done with 200 samples per iteration.

with a degenerate ground model is presented. The underlying graph has ten nodes and two max-cliques of size $K = 5$. One can observe from the charts for both models VBS is able to outperform the trainable GBS approach of Banchi et al. (2020). The added variability is primarily due to sampling at each step. The figures also present the sampling variation (p=0.1) during different trials in the run. It is clear that the GBS curve is statistically better than the VBS curve.

Next, the experiment is repeated with the aforementioned families of random graphs. These results are presented in the Appendix (Figure 7). The first row presents results on instances of random Barabasi-Albert graphs. These graphs have many cliques of sizes three and four, leading to multiple local optimas. The second row illustrates the result of training a VBS model on random Erdos-Renyi graphs with ten vertices. We can observe that both GBS and VBS can with high probability ($> 80 - 90\%$) sample the energy minimum. However it is also clear that the VBS trained model can sample the configuration corresponding to largest clique in the graph with a higher success rate than the GBS approach. The performance curves also make it clear from these results that the VBS behaviour performance is fundamentally distinct from the GBS one (e.g. see subfigures 2,5).

4.2 Generative Modelling/Image Generation

Next we use the VBS scheme to learn a simple generative model. For this experiment, we used the test set of the UCI digits dataset (Alpaydin and Kaynak, 1998)[1].

[1] Available at https://archive.ics.uci.edu/ml/datasets/Optical+Recognition+of+Handwritten+Digits

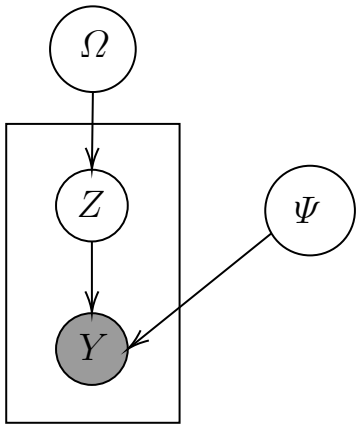

Fig. 4: Generative Model for Digit images

The UCI DIGITS dataset consists of 1797 data observations. Each observation is an 8×8 image of a handwritten digit. Since the simulation of a 64 photons BS scheme is not feasible on classical machines, we use a low dimensional embedding approach. The VBS model sample vectors from a 5-dimensional latent space which then constructs the images through a classical conditional generator. The complete generative model is depicted in Figure 9. In the figure, Y is a random variable that denotes the training images, Z is a latent variable, and P_Ω is the distribution of Z parameterized by Ω. The parameters Ψ are the parameters of the conditional generator. z and y are used to denote samples of the variables Z and Y respectively. We follow the same experiment design on 5000 randomly selected samples from MNIST.

In our experiments, we model the conditional distribution of the output Y given the latent z as a Gaussian variable. The distribution of the latent variable z is given by the output distribution of the Boson sampler. The combined log-likelihood of an observation y is then given by:

$$L(y; W, b, \Omega) = \int N(y; W\bar{z} + b, \Sigma) P_{\text{VBS}}(z; \Omega) dz$$

Here Ω refers to the parameters of the VBS scheme (i.e $\Omega = \{\boldsymbol{\theta}, \boldsymbol{\phi}, \boldsymbol{\varphi}\}$ P_{VBS} is the induced distribution over z by Ω. $\bar{z}$ refers to the normalized value of z. Since z itself is a discrete distribution with the number of bosons in each mode being a non-negative integer, we scale the output by its norm to make it approximately continuous. This, in turn, determines the mean vector of the Gaussian distribution N. Σ is assumed to be a diagonal matrix, and W, b are parameters learned on a classical device. Hence in this case $\Psi = \{\boldsymbol{W}, \boldsymbol{b}, \boldsymbol{\Sigma}\}$ We assume that all observed images are independent draws from the generating distribution. The total likelihood of the data is then just the product of likelihoods for each observation.

We train the model to maximize the log-likelihood of the data with a variant of the EM algorithm (Dempster et al., 1977). In the EM algorithm, each iteration consists of repeated application of the E-step and the M-step. In the E-step, the data log-likelihood conditioned on the observed variables is computed. On the other hand, in the M-step, the likelihood obtained in the E-step is maximized with respect to the model parameters. The Monte-Carlo Expectation Maximization

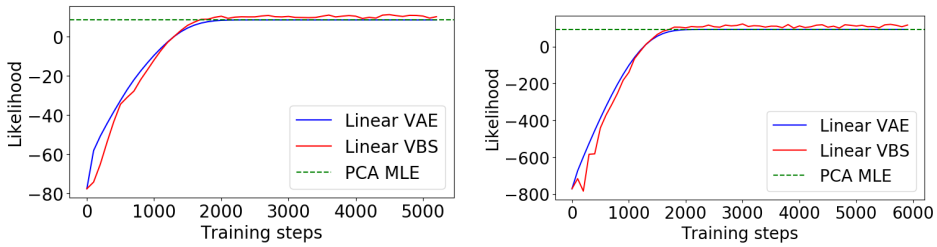

Fig. 5: Test likelihood across different training iterations trained generative model described in Section 4.2 with 5 dimensional latent space. Left (a) presents the number on DIGITS dataset, while (b) shows results on a subset of MNIST

(MCEM) algorithm (Wei and Tanner, 1990) is a variant of the classic EM; often used for high-dimensional data or when the integral required in the E-step is intractable. The key difference between the two is that the MCEM uses a Monte-Carlo approximation to the conditional expectation during the E-step. Since computing the exact output distribution of the BS scheme is generally intractable, while quantum devices can sample from it easily, MCEM is a better choice for training such models. Specifically the Q function for the MCEM algorithm in our case is given by:

$$Q(\Omega_t|\Omega_{t-1}, y) = \mathbb{E}_{z\sim P_{\text{VBS}}(\Omega_{t-1})} \log N(y; W\bar{z} + b, \Sigma) \tag{5}$$

$$\approx \sum_{z_i \sim P_{\text{VBS}}(\Omega_{t-1})} \log N(y; W\bar{z_i} + b, \Sigma) \tag{6}$$

where the sum is over samples z_i drawn from P_{VBS}. This Q function is then optimized by gradient descent to estimate Ω. Similar to REINFORCE, monte-carlo EM adds extra variability and generally requires lower learning rate for smooth learning. However these problems are not as significant for these experiments as the likelihood loss is dominated by the decoder terms instead of the prior.

We also compare our results against a VAE (Kingma and Welling, 2019) with the same sized latent space and a linear decoder like in the VBS model. A linear decoder was chosen as such a model with Gaussian prior provides an exactly computable likelihood. One can use more complex decoder models for better sample quality, but our goal in this work is simply a working BS based generative model. Both models were trained with Adam optimizer with a learning rate of 5e-4.

Note that under a Gaussian prior the given generative model corresponds exactly to the PCA decomposition of the data, which can be analytically computed. We present this exact likelihood (pPCA MLE value) in Figure 5. Furthermore given the low latent dimensionality, a generic prior and the significantly higher contribution from the conditional model, the learnt VBS based distribution in this case is expected to be similar to the PCA decomposition. This can be seen from

Figure 5, where we plot the data log-likelihood as learnt by the PCA, VAE and VBS models. We can also observe that the VBS model improves the likelihood by around 3%

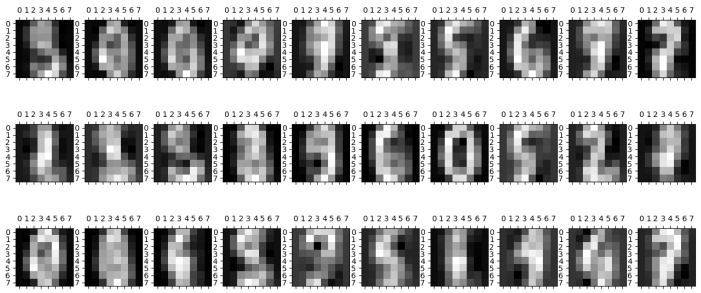

(a)

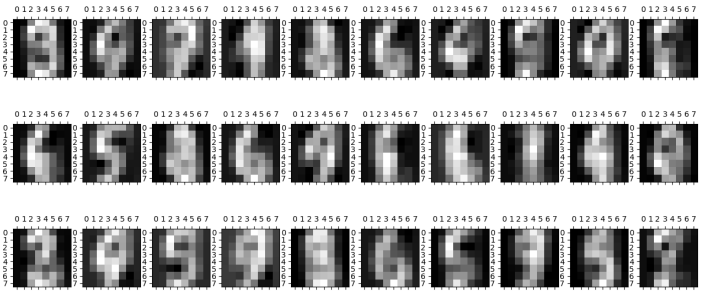

(b)

Fig. 6: Samples generated by the trained generative model described in Section 4.2 with 5 dimensional latent space. Upper rows (a) presents samples from the VBS model and lower rows (b) shows sample from a linear VAE model.

The samples from both VAE and VBS runs are presented in Figure 6. The upper rows present samples from the VBS model trained with MCEM. On the lower rows we have samples from the baseline linear VAE. It is evident from the image quality that a 5-dimensional latent space is too small. However this is primarily a function of the decoder. For example, MNIST has high quality reconstructions from a two dimensional latent space with a non-linear decoder [2].

[2] `https://github.com/lttsh/VariationalAutoEncoder-MNIST`

On the other hand a linear decoder with (a gaussian conditional generator) is equivalent to PCA and produces similarly diffused images. However, even with the blurred images, the digit-like structure of these sampler is clear. A qualitative examination shows that both the VBS and the VAE models produces samples of a similar nature. Combined with the likelihood results, this experiment provides evidence that the VBS model is at least as powerful as a VAE in this setting.

5 Conclusion

In this paper, we demonstrate that boson sampling can be used to solve practical problems in machine learning and optimization. Towards this goal, we developed a hybrid quantum-classical variational scheme labeled Variational Boson Sampling (VBS). VBS introduces an optimizable parametric structure into the boson evolution operation and uses that as a variational ansatz. We then experiment with this framework for training VBS distributions for problems in ising optimization and generative modelling.

For optimization, we tried an algorithm where the VBS is used to generate samples that can be mapped to the states of an Ising model. We then use REINFORCE (Williams, 1992) to get stochastic gradients of the parameters device in order to maximize the probability of sampling the ground state of the Ising model. In generative modelling, we show that a VBS-based scheme trained using an EM algorithm is competitive with a VAE (Kingma and Welling, 2019) of similar capacity. While sampling based methods do add extra variability in learning, we do not believe this to be a major issue as physical implementation of a VBS is extremely time-efficient in producing samples. For example current quantum devices can draw 10^5 samples per second (Vaidya et al., 2020). As such neither variability nor the total runtime for training VBS scheme is likely to be high.

Our results have shown that the VBS scheme can be used to implement algorithms for practical problems; we hope this sparks more research into variational boson sampling ansatz in the future. One future research direction is formalizing cases when a VBS-based scheme outperforms alternative algorithms such as VQE (Wang et al., 2019) or QAOA (Farhi et al., 2014) for standard qubit based devices. Another potential research direction is to develop schemes to approximate the gradients of the scheme using a quantum device. Finally recent work (Ostaszewski et al., 2021) has used deep Q-learning to optimize larger quantum circuits, and using such techniques can be of potential use in improving training for boson sampling as well.

References

S. Aaronson and A. Arkhipov. The computational complexity of linear optics. In *43rd Annual ACM Symp. Theory of Computing*, pages 333–342, 2011.

E. Alpaydin and C. Kaynak. Cascaded classifiers. *Kybernetika*, 34:369–374, 1998.

J. M. Arrazola and T. R. Bromley. Using gaussian boson sampling to find dense subgraphs. *Physical review letters*, 121(3):030503, 2018.

F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.

L. Banchi, N. Quesada, and J. M. Arrazola. Training gaussian boson sampling distributions. *Physical Review A*, 102(1):012417, 2020.

A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani. On the complexity and verification of quantum random circuit sampling. *Nature Physics*, 15 (2):159–163, 2018. ISSN 1745-2481. URL `http://dx.doi.org/10.1038/s41567-018-0318-2`.

M. J. Bremner, R. Jozsa, and D. J. Shepherd. Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 467 (2126):459–472, 2011.

Bromley et al. Applications of near-term photonic quantum computers. *Quantum Science and Technology*, 2021.

A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1):1–22, 1977.

E. Farhi, J. Goldstone, and S. Gutmann. A quantum approximate optimization algorithm. *arXiv preprint arXiv:1411.4028*, 2014.

Y. Feng, R. Shi, J. Shi, W. Zhao, Y. Lu, and Y. Tang. Arbitrated quantum signature protocol with boson sampling-based random unitary encryption. *Journal of Physics A: Mathematical and Theoretical*, 53(13):135301, 2020.

Z. J. Gao, N. Pansare, and C. Jermaine. Declarative parameterizations of user-defined functions for large-scale machine learning and optimization. *IEEE Transactions on Knowledge and Data Engineering*, 31(11):2079–2092, 2018.

C. Gogolin, M. Kliesch, L. Aolita, and J. Eisert. Boson-sampling in the light of sample complexity. *arXiv preprint arXiv:1306.3995*, 2013.

L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, 1996.

G. G. Guerreschi. Boson sampling for molecular vibronic spectra. *Bulletin of the American Physical Society*, 60, 2015.

C. S. Hamilton, R. Kruse, L. Sansoni, S. Barkhofen, C. Silberhorn, and I. Jex. Gaussian boson sampling. In *2017 Conference on Lasers and Electro-Optics (CLEO)*, 2017.

Z. Huang, P. P. Rohde, D. W. Berry, P. Kok, J. P. Dowling, and C. Lupo. Boson sampling private-key quantum cryptography. 2019.

D. P. Kingma and M. Welling. An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning*, 12(4):307–392, 2019. ISSN 1935-8245.

D. Liu, S.-J. Ran, P. Wittek, C. Peng, R. B. García, G. Su, and M. Lewenstein. Machine learning by unitary tensor network of hierarchical tree structure. *New Journal of Physics*, 21(7):073059, 2019.

J.-G. Liu and L. Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6):062324, 2018a.

J.-G. Liu and L. Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6):062324, 2018b.

A. Lucas. Ising formulations of many np problems. *Frontiers in Physics*, 2, 2014. ISSN 2296-424X. https://doi.org/10.3389/fphy.2014.00005. URL http://dx.doi.org/10.3389/fphy.2014.00005.

A. P. Lund, A. Laing, S. Rahimi-Keshari, T. Rudolph, J. L. O'Brien, and T. C. Ralph. Boson sampling from a gaussian state. *Physical review letters*, 113(10): 100502, 2014.

A. P. Lund, M. J. Bremner, and T. C. Ralph. Quantum sampling problems, bosonsampling and quantum supremacy. *npj Quantum Information*, 3(1), 2017. ISSN 2056-6387. URL http://dx.doi.org/10.1038/s41534-017-0018-2.

J. R. McClean, J. Romero, R. Babbush, and A. Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18 (2):023023, 2016.

N. Moll, P. Barkoutsos, L. S. Bishop, J. M. Chow, A. Cross, D. J. Egger, S. Filipp, A. Fuhrer, J. M. Gambetta, M. Ganzhorn, et al. Quantum optimization using variational algorithms on near-term quantum devices. *Quantum Science and Technology*, 3(3):030503, 2018.

M. E. Morales, T. Tlyachev, and J. Biamonte. Variational learning of grover's quantum search algorithm. *Physical Review A*, 98(6):062333, 2018.

T. Morimae, H. Nishimura, K. Fujii, and S. Tamate. Classical simulation of dqc12 or dqc21 implies collapse of the polynomial hierarchy. *Phys. Rev. Lett*, 81:5672, 1998.

G. Muraleedharan, A. Miyake, and I. H. Deutsch. Quantum computational supremacy in the sampling of bosonic random walkers on a one-dimensional lattice. *New Journal of Physics*, 21(5):055003, 2019. URL https://doi.org/10.1088/1367-2630/ab0610.

M. A. Nielsen and I. Chuang. Quantum computation and quantum information, 2002.

C. Oh et al. Classical simulation of lossy boson sampling using matrix product operators. *Physical Review A*, 2021.

M. Ostaszewski, L. M. Trenkwalder, W. Masarczyk, E. Scerri, and V. Dunjko. Reinforcement learning for optimization of variational quantum circuit architectures. *arXiv preprint arXiv:2103.16089*, 2021.

J. Otterbach, R. Manenti, N. Alidoust, A. Bestwick, M. Block, B. Bloom, S. Caldwell, N. Didier, E. S. Fried, S. Hong, et al. Unsupervised machine learning on a hybrid quantum computer. *arXiv preprint arXiv:1712.05771*, 2017.

M. Schuld, K. Brádler, R. Israel, D. Su, and B. Gupt. Measuring the similarity of graphs with a gaussian boson sampler. *Physical Review A*, 101(3), 2020. ISSN 2469-9934. URL http://dx.doi.org/10.1103/PhysRevA.101.032314.

S. Shankar and D. Towsley. Bosonic random walk networks for graph learning. *Quantum Tensor Networks in Machine Learning*, 2020.

P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2):303–332, 1999.

Vaidya et al. Broadband quadrature-squeezed vacum and nonclassical photon number correlations on nanophotonic device. *Science Advances*, 2020.

L. Valiant. The complexity of computing the permanent. *Theoretical Computer Science*, 8(2):189–201, 1979. https://doi.org/10.1016/0304-3975(79)90044-6.

D. Wang, O. Higgott, and S. Brierley. Accelerated variational quantum eigensolver. *Physical review letters*, 122(14):140504, 2019.

G. C. Wei and M. A. Tanner. A monte carlo implementation of the em algorithm and the poor man's data augmentation algorithms. *Journal of the American statistical Association*, 85(411):699–704, 1990.

R. J. Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.

C.-H. Yu, F. Gao, Q.-L. Wang, and Q.-Y. Wen. Quantum algorithm for association rules mining. *Physical Review A*, 94(4):042311, 2016.

C.-H. Yu, F. Gao, and Q. Wen. An improved quantum algorithm for ridge regression. *IEEE Transactions on Knowledge and Data Engineering*, 2019.

J. Zeng, Y. Wu, J.-G. Liu, L. Wang, and J. Hu. Learning and inference on generative adversarial quantum circuits. *Physical Review A*, 99(5):052306, 2019.

H.-S. Zhong, H. Wang, Y.-H. Deng, M.-C. Chen, L.-C. Peng, Y.-H. Luo, J. Qin, D. Wu, X. Ding, Y. Hu, et al. Quantum computational advantage using photons. *Science*, 370(6523):1460–1463, 2020.

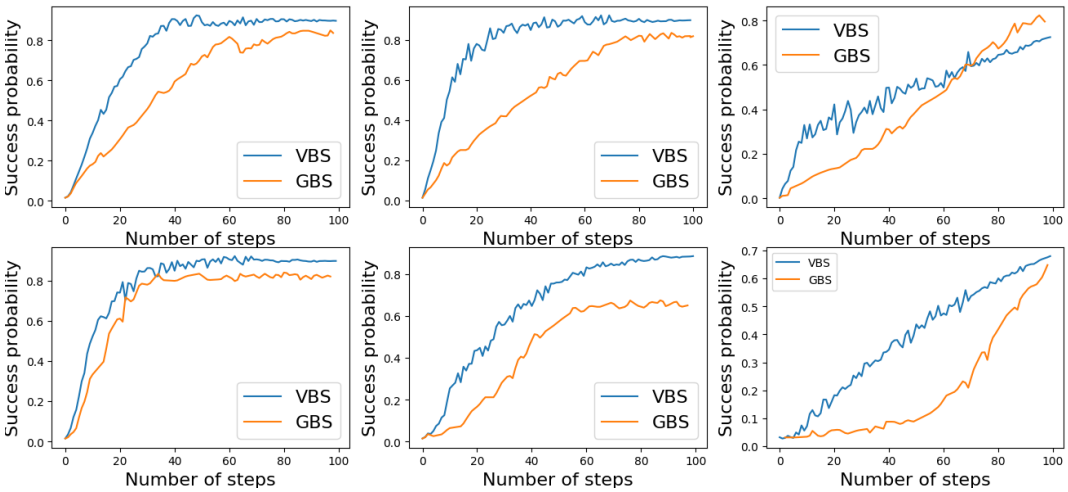

Fig. 7: Success rate over time as the training progresses, as in Figure 3, Graphs (a),(b),(c) are random Barabasi-Albert graphs with ten vertices, built starting from a clique of five vertices and attaching new vertices, each connected to three random nodes. Graphs (d),(e),(f) are random Erdos-Renyi graphs with ten vertices and probability $p = 0.5$ of adding an edge between pairs of vertices.

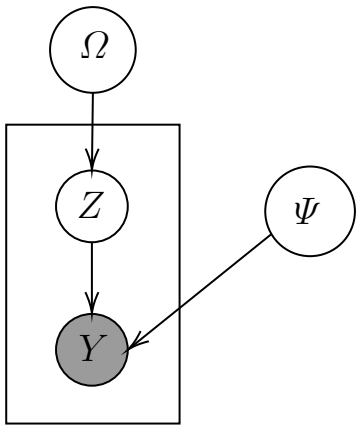

Fig. 9: Generative Model for Digit images

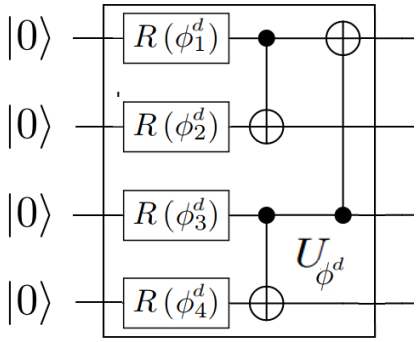

(a)

Fig. 8: Variational QCBM circuit architecture from Liu and Wang (2018b) with parameters ϕ^d