# UNIVERSAL DISCRIMINATIVE QUANTUM NEURAL NETWORKS

**Anonymous authors**
**Paper under double-blind review**

## ABSTRACT

Quantum mechanics fundamentally forbids deterministic discrimination of quantum states and 9b8d to discriminate among various pure and mixed quantum states exhibiting a trade-off between minimizing erroneous and inconclusive outcomes with comparable performance to theoretically optimal POVMs. We train the circuit on different classes of quantum data and evaluate the generalization error on unseen mixed quantum states. This generalization power hence distinguishes our work from standard circuit optimization and provides an example of quantum machine learning for a task that has inherently no classical analog.

## 1 INTRODUCTION

The interface of quantum physics and machine learning has recently received a considerable amount of interest. Two complementary methodologies have been developed to address the question whether quantum mechanics can help with solving machine learning tasks or similarly whether machine learning techniques could help to solve problems in quantum computation and many-body condensed matter systems more efficiently (Ciliberto et al. (2018); Biamonte et al. (2017); Schuld et al. (2015)). Using the circuit model of universal fault-tolerant quantum computers, it has been shown that quantum computation offers a range of machine learning algorithms (Lloyd et al. (2014; 2013); Rebentrost et al. (2014); Wiebe et al. (2015)) which could lead to considerable speed-ups under certain assumptions over classical counterparts. The underlying property which results in such quantum advantage is the ability of quantum computers to perform certain linear algebra operations faster than classical machines (Harrow et al. (2009); Wossnig et al. (2017); Childs et al. (2017); Ambainis (2010); Dervovic et al. (2018)). The algorithmic complexity for some of these quantum learning schemes is in principle $O(\text{polylog}(N))$ for input dimension $N$ instead of the classical $O(\text{poly}(N))$ which can manifest itself as exponential speedup when applied to quantum data (Lloyd et al. (2014)). However, these algorithms have been shown to come with certain caveats when applied to classical data. The most relevant one is the preparation of quantum states that encode classical information and is believed to have in the worst case a polynomial scaling in the input dimension of the data (Aaronson (2015); Childs (2009)), diminishing their advantages in the first place.

Here we take a different approach and utilize a hybrid quantum-classical setup to directly learn the design of shallow quantum circuits for an inherently quantum-mechanical task, which has no classical counterpart based on the non-orthogonal nature of quantum states. Hence we solve a problem where a comparison with classical algorithms is not appropriate unlike for the case of the above-mentioned algorithms. Recently a string of works focusing on using quantum-classical hybrid methods for training quantum circuits for a wide range of tasks have been proposed (Banchi et al. (2016); Wan et al. (2017); Innocenti et al. (2018); Romero et al. (2017); Mitarai et al. (2018); Farhi & Neven (2018); Verdon et al. (2017); Li & Benjamin (2017); Grant et al. (2018); Schuld et al. (2018)).

Our theoretical foundation is based on an important primitive know as *quantum state discrimination*. Quantum state discrimination aims to identify the state of a quantum system among a set of possible

*Email: l.wossnig@cs.ucl.ac.uk
†Email: mohseni@google.com

candidates known *a priori*. A key feature for discrimination of quantum states is that a deterministic discrimination is impossible when the complex vectors representing the candidate states are not orthogonal, i.e. when their overlaps are non-zero. The central task of quantum state discrimination is to find the optimal probabilistic discrimination for non-orthogonal quantum states (Barnett & Croke (2009); Bergou (2007)).

In this work, we show that a universal quantum circuit can be trained as a discriminator for classification of non-orthogonal quantum data sampled from various different probability distributions. Our discriminator can achieve a near-zero error rate by producing inconclusive signals. Notably, unlike previous works on quantum state discrimination, we focus on the generalization ability of our circuit, i.e., we train the circuit on a specific range of the parameters with the goal of maximizing its generalization performance, hence considering a learning framework. This distinguishes our work from the pure optimization problem for the state discrimination task, which is optimizing the circuit to distinguish only a concrete set of states.

Our circuit topologies comprise gates from a universal gate set consisting of the C-NOT gate and single-qubit gates, motivated by the fact that implementations of these gates are known for the currently mostly used experimental architectures. Furthermore, our decompositions are nearly optimal in terms of the number of C-NOT gates which is an important feature for an implementation on near-term devices. The quantum circuits we apply here can be viewed as a form of quantum neural networks with non-unitary layers, i.e., the generalized measurements, leading to sufficient non-linearities. The reason is the close mathematical relation between the unitary circuit structures and classical unitary-weight neural network structures (Schuld et al. (2018); Arjovsky et al. (2016); Hyland & Rätsch (2017)). It has been shown that the unitarity of the layers indeed is optimal since it can avoid the exploding or vanishing gradient problem if the activation function is chosen adequately, e.g., a ReLu (Arjovsky et al. (2016)). Since quantum circuits do not include any activation functions between the layers, unless a projective measurement is performed, this advantage is immediately transferable to the case of quantum circuit training. As pointed out already in Schuld et al. (2018), it has further been shown that unitary weight matrices allow for gradient descent methods to converge independently of the circuit depth (Saxe et al. (2013)), which is important for the training of circuits with larger circuit depth. A high-level description of a comparison of our quantum circuit learning and an analogy with a neural network structure is given in Fig 1. Recently, unitary architectures have also been considered for the training of classical neural networks (Arjovsky et al. (2016); Hyland & Rätsch (2017)).

The discriminative quantum neural networks introduced here could have a broad range of applications. Quantum state discrimination by itself plays a key role in quantum information processing protocols and is used in quantum cryptography (Bennett (1992)), quantum cloning (Duan & Guo (1998)), quantum state separation and entanglement concentration (Chefles (2000)). Our shallow quantum circuit learning could further be used to construct quantum repeaters and state purification units within quantum communication networks. It could also have a wide range of applications in quantum meteorology (Giovannetti et al. (2011)) quantum sensing (Degen et al. (2017)), quantum illumination (Lloyd (2008)), and quantum imaging Chen et al. (2012) as a systematic way of engineering structured receivers. In general, our discriminative networks could be used as a general quantum circuit verification units to examine the outputs of other quantum circuits with possible applications to quantum versions of Generative Adversarial Networks (GANs) (Goodfellow et al. (2014)).

## 2 Dataset and our approach

We focused on the classification of several families of non-orthogonal quantum states, each coming from a source of stochastic nature:

$$\rho_i(a_i),\ a_i \sim \alpha_i, \tag{2.1}$$

where $i$ labels different classes, the $a_i$ denote the parameters for the density matrices and are determined by a stochastic source. The stochastic feature of the source is modeled as a probability distribution $\alpha_i$ and $a_i$ are samples drawn according to the distribution $\alpha_i$ when training a quantum circuit to classify the quantum data.

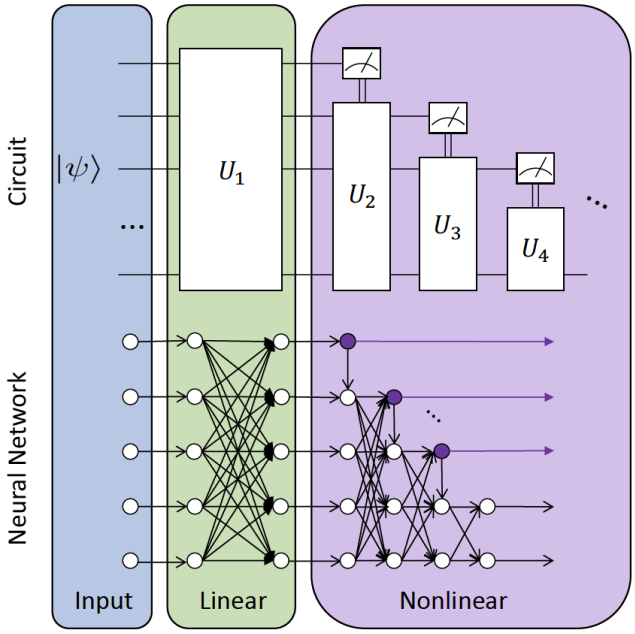


Figure 1: Universal quantum circuit learning construction can be viewed as quantum neural networks. The unitary evolution of quantum states could be considered as a layer of symmetric fully connected neurons. Multiple layers of generalized quantum measurements (POVMs) brings in sufficient non-linearities for performing classification or discrimination tasks.

For this work, we restricted our attention to the classification of two families of quantum states stored in 2 qubits. Our first family consists of pure states, parametrized by a real number $a \in [0, 1]$:

$$\psi_1(a) = \left(\sqrt{1-a^2}, 0, a, 0\right), \; \rho_1 = |\psi_1(a)\rangle\langle\psi_1(a)| \,. \tag{2.2}$$

The second family consists of mixed states $\rho_2(b)$ where $b \in [0, 1]$. Specifically,

$$\psi_{2/3} = \left(0, \pm\sqrt{1-b^2}, b, 0\right), \; \rho_2 = \frac{1}{2}|\psi_2\rangle\langle\psi_2| + \frac{1}{2}|\psi_3\rangle\langle\psi_3| \,. \tag{2.3}$$

The overlap between $\psi_1$ and $\psi_{2/3}$ is $ab$, and hence our data in two families are non-orthogonal. In the case when $a$ is a fixed value, and $b = \frac{1}{\sqrt{2}}$, (i.e. when they both follow the Dirac delta distribution), the maximal success rate for unambiguously discriminating between $\rho_1$ and $\rho_2$ has been theoretically studied, and an experimental demonstration is available (Mohseni et al. (2004)). The specific distribution we have tested in our experiments are summarized in Table 1. To generate the data for the training, validation, and testing of our circuits, we randomly sampled points from the corresponding distributions.

**Approach.** Generally, there are two different strategies to cope with our inherent inability to perform deterministic discrimination: (a) *Minimum-error discrimination*: In this strategy, the task is to minimize the probability that the inevitable errors occur in the classification. (b) *Unambiguous discrimination*: In this strategy, the discriminator has one more output prediction than the number of classes it tries to classify: the inconclusive outcome. The task is to eliminate the error rate of the discriminator while minimizing the probability of the inevitable inconclusive outcomes. A pure unambiguous discrimination with strictly 0 error rate is not guaranteed to be possible. From the perspective of numerical optimization, one needs to allow for some small but non-zero errors to happen.

| | Family 1 ($\rho_1(a)$) | Family 2 ($\rho_2(b)$) |
|---|---|---|
| Case 1 | $a \in [0,1]$ | $b \in [0,1]$ |
| Case 2 | $a \approx 0.25$ | $b \approx \frac{1}{\sqrt{2}}$ |
| Case 3 | $a \approx 0.25$ | $b \in [0,1]$ |
| Case 4 | $a \in [0,1]$ | $b \approx \frac{1}{\sqrt{2}}$ |

Table 1: A summary of different test cases we classified in this work. Here $a(b) \in [0,1]$ represents that $a(b)$ follows a uniform distribution in $[0,1]$. $a(b) \approx 0.25(\frac{1}{\sqrt{2}})$ represents that $a(b)$ follows a normal distribution with mean $0.25(\frac{1}{\sqrt{2}})$, standard deviation $0.05$, and which is truncated in $[0,1]$.

In this work, we used the machine learning approach to train a universal quantum circuit capable of giving any allowed quantum measurements with four possible measurement outcomes $m_{i_2 i_1}$, where $i_1, i_2 \in \{0,1\}$ are the measurement outcomes of the first and the second qubit respectively. The decomposition of this circuit is discussed in Appendix A. By assigning $m_{00}$ or $m_{10}$ to the input $\rho_1(a)$, $m_{01}$ to the input $\rho_{2/3}$, and $m_{11}$ to the inconclusive measurement, this circuit acts as a discriminator with the various probabilities (success probability $P_{\text{suc}}$, error probability $P_{\text{err}}$ and inconclusive probability $P_{\text{inc}}$) with respect to the input data naturally defined. We trained this circuit both as a minimal-error discriminator and an unambiguous discriminator by iteratively measuring the three probabilities and modifying the parameters of the circuit accordingly using the Adam optimization algorithm (Kingma & Ba (2014)). In this work, we performed experiments on simulated quantum computers, where the probabilities were available and the gradients were calculated by the forward difference formula. We note that on real quantum computers, these probabilities can be inferred through repeated measurements on replicated data, up to some precision.

We used the loss function defined in Eq. 2.4. There we assumed that for each family of quantum states, we had been supplied with a set $S_i$ of training samples, where each class was labeled by $i$. Then $|S_i|$ denotes the number of samples in the training sets $S_i$, $\alpha_{\text{err}}$ is the penalty for making errors, and $\alpha_{\text{inc}}$ is the penalty for giving inconclusive outcomes. $P_{\text{suc}}(\psi)/P_{\text{err}}(\psi)/P_{\text{inc}}(\psi)$ are the probabilities of giving a correct/erroneous/inconclusive measurement outcome for the specific input quantum data $\rho_i$. This loss function measures the performance of our quantum circuit as a minimal-error discriminator (when $\alpha_{\text{err}} < \alpha_{\text{inc}}$) or as an unambiguous discriminator (when $\alpha_{\text{err}} > \alpha_{\text{inc}}$).

$$
\begin{aligned}
J = &\sum_i \frac{1}{|S_i|} \sum_{a_i \in S_i} (1 - P_{\text{suc}}(\rho_i(a_i))) \\
&+ \alpha_{\text{err}} \sum_i \frac{1}{|S_i|} \sum_{a_i \in S_i} P_{\text{err}}(\rho_i(a_i)) \\
&+ \alpha_{\text{inc}} \sum_i \frac{1}{|S_i|} \sum_{a_i \in S_i} P_{\text{inc}}(\rho_i(a_i)).
\end{aligned} \tag{2.4}
$$

For our specific problem of classifying $\rho_1$ and $\rho_2$ as defined in Eq. 2.2 and Eq. 2.3, we defined the success/erroneous/inconclusive rates in Equation 2.5 to summarize and compare the performance off different trainings:

$$
\begin{aligned}
P_s &= \frac{1}{3} P_s(\rho_1)_{\text{avg}} + \frac{2}{3} P_s(\rho_2)_{\text{avg}} \\
&= \frac{1}{3} P_s(\psi_1)_{\text{avg}} + \frac{1}{3} P_s(\psi_2)_{\text{avg}} + \frac{1}{3} P_s(\psi_3)_{\text{avg}}
\end{aligned} \tag{2.5}
$$

where $s$ is a place-holder for *suc* (successful), *err* (erroneous) or *inc* (inconclusive). The subscript *avg* means that the probabilities are calculated as the value averaged in all samples available in either the training set, or the test set (but not both). The choice of weights ($\frac{1}{3}$and $\frac{2}{3}$) in the Equation 2.5 was made to be consistent with the paper Mohseni et al. (2004).

## 3 THEORETICAL ANALYSIS

Here we describe a theoretical result which we used as a reference for our numerical results. Assuming that we have a quantum measurement described by a positive-operator valued measure with elements $\{\Pi_i\}_{i\in\mathbb{N}}$ then, we can define the average probability of $i$, i.e., an occurance of a measurement outcome $i$ as

$$\begin{aligned}\int \mathrm{Tr}(\Pi_i\rho(a))\mathrm{d}a &= \mathrm{Tr}\left[\int \Pi_i\rho(a)\mathrm{d}a\right] = \mathrm{Tr}\left[\Pi_i \int \rho(a)\mathrm{d}a\right] \\ &= \mathrm{Tr}[\Pi_i\rho], \end{aligned} \tag{3.1}$$

where $\rho \equiv \int \rho(a)\mathrm{d}a$, and the integration of the matrix is done in an element-wise fashion. Therefore, if $\mathrm{Tr}(\Pi_i\rho) = 0$ for some $i$, then $\int_D \mathrm{Tr}(\Pi_i\rho(a)) = 0$ for any subset $D$ with non-zero measure in the whole parameter space of $a$. This is due to the fact that $\mathrm{Tr}[\Pi_i\rho(a)] \geq 0$ for any parameter $a$.

Applied to our specific cases, the problem of unambiguous discrimination of $\rho_1$ and $\rho_2$ is equivalent to the problem of unambiguous discrimination of the family $\rho_1(a), \forall a$, from the family $\rho_2(b), \forall b$. Here $\rho_1 = \int_a \rho_1(a)P_1(a)\mathrm{d}a$, $\rho_2 = \int_a \rho_2(a)P_2(a)\mathrm{d}b$, and $P_1(a)$ and $P_2(b)$ are the probability distribution of $a$ and $b$ respectively. Let $\{\Pi_1, \Pi_2, \Pi_?\}$ be a POVM that unambiguously classifies all members of the two families $\rho_1(a)$ and $\rho_2(b)$, for all possible parameters. That is when $\Pi_?$ corresponds to the inconclusive outcome and

$$\mathrm{Tr}(\Pi_2\rho_1(a)) = 0,\ \forall a \in D_a,$$
$$\mathrm{Tr}(\Pi_1\rho_2(b)) = 0,\ \forall b \in D_b,$$

then $\mathrm{Tr}(\Pi_1\rho_2) = \mathrm{Tr}(\Pi_2\rho_1) = 0$, and vice versa. Using this formalism, we theoretically analyzed the different cases we described in Table 1 based on the works of Raynal et al. (2003) and Barnett & Croke (2009), and the results are displayed in Table 2. Note that these are average case success probabilities.

| | Family 1 ($\rho_1$) | Family 2 ($\rho_2$) | $P_{\text{suc}}$ |
|---|---|---|---|
| Case 1 | $a \in [0,1]$ | $b \in [0,1]$ | 0.67 |
| Case 2 | $a \approx 0.25$ | $b \approx \frac{1}{\sqrt{2}}$ | 0.64 |
| Case 3 | $a \approx 0.25$ | $b \in [0,1]$ | 0.76 |
| Case 4 | $a \in [0,1]$ | $b \approx \frac{1}{\sqrt{2}}$ | 0.55 |

Table 2: A summary of maximal success rate when the error rate is exactly 0 for different test cases we have classified in this work.

## 4 NUMERICAL RESULTS ON SIMULATED QUANTUM COMPUTERS

We presented the results of training the discriminator to be an unambiguous discriminator for different distributions summarized in Table 1. To balance between eliminating the error rate ($P_{\text{err}}$) while minimizing the inconclusive rate ($P_{\text{inc}}$), we used a specific training strategy described in the following. We first prioritized a smaller inconclusive rate by starting with a zero penalty for error ($\alpha_{\text{inc}} > \alpha_{\text{err}} = 0$) and then increased the $\alpha_{\text{err}}$ in a step-wise manner until a certain objective error rate is achieved. Similar optimization procedures have been used in the context of variational auto-encoders both in classical machine learning (Snderby et al. (2016)), and in quantum machine learning applications (Rocchetto et al. (2017)) and are typically referred to as warm-up scheme. Using such, we trained our circuits to unambiguously discriminate the two families of quantum states, and observed a convergence towards the theoretical success rates (Section 3) with increasing amount of data used to train the discriminator. Notably, we did not observe any signs of overfitting despite the increasing dataset size (Figure 2 a).

### 4.1 TRADE-OFF BETWEEN THE ERROR RATE AND THE INCONCLUSIVE RATE

In another test, we showed that our model is able to obtain a much higher success rate ($P_{\text{suc}}$) if the inconclusive rate may be lowered, i.e., if we allow for a slightly higher error rate during the training

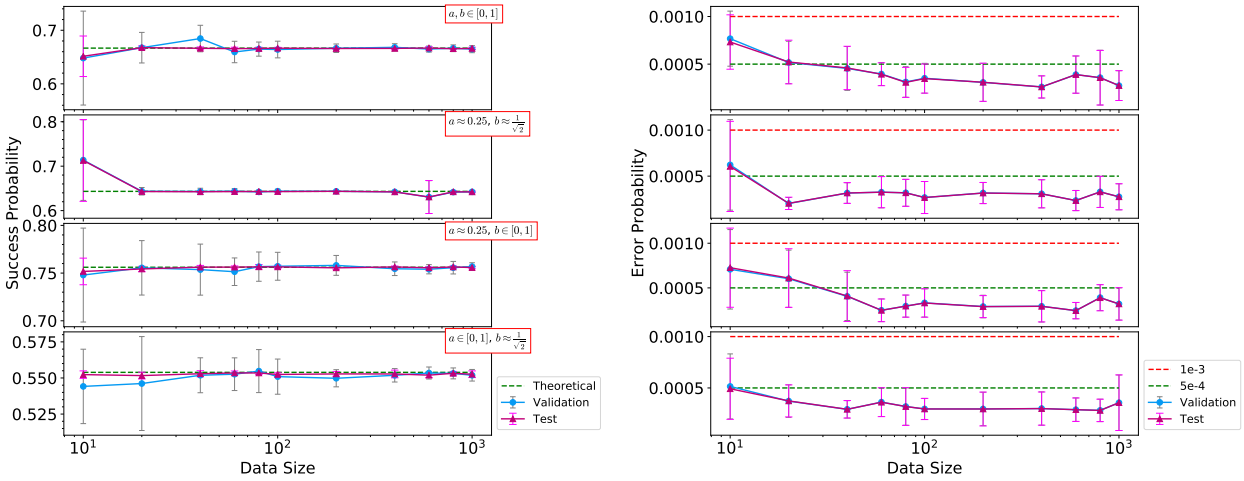


(a) The success rates converge to the theoretical result (the green dashed line in the figure) as the training data size increased. The performance on the testing (purple line) and the validation dataset also converges with the growing data size. The abnormal deviation at point the data size 600 for the case $a \approx 0.25$ and $b \approx 1/\sqrt{2}$ was due to unsuccessful optimization.

(b) Most of the error rate for the corresponding trainings are well below 0.0005, indicating an unambiguous discriminator was achieved.

Figure 2: Unambiguous classification of non-orthogonal quantum data sampled from different probability distributions. The data were averaged over 10 repeated trails starting with random initializations and the bars indicate the standard deviations. The training, validation and test data were sampled from the corresponding distributions. The data size indicates the size of the training and the validation set. The test data was fixed at a size of $10^4$ points for each family and each distribution.

compared to the previous results. This hints at a trade-off between the error rate ($P_{err}$) and the inconclusive rate ($P_{inc}$) which can be utilized in real-world applications.

For the dataset "Case 4" in Table 1 we fixed the two penalties ($\alpha_{\text{err}}$ and $\alpha_{\text{inc}}$) during then training and observed a gradual transition from unambiguous-like classification (characterized by a near-zero error probability) to minimal-error-like classification (characterized by the near-zero inconclusiveness) where we used varying penalties (Figure 3 a-c) throughout the training. The gain from allowing a small but non-zero error rate allowed much higher success rates which are unpredicted by the theory. We note that introducing the penalty made the training process also more stable (Figure 3 a). Our results showed that in comparison with the case without penalties ($\alpha_{\text{err}} = \alpha_{\text{inc}} = 0$), their introduction acted as a form of regularization and could be adjusted to give a higher success probability or a lower inconclusiveness rate for the final model (Figure 3).

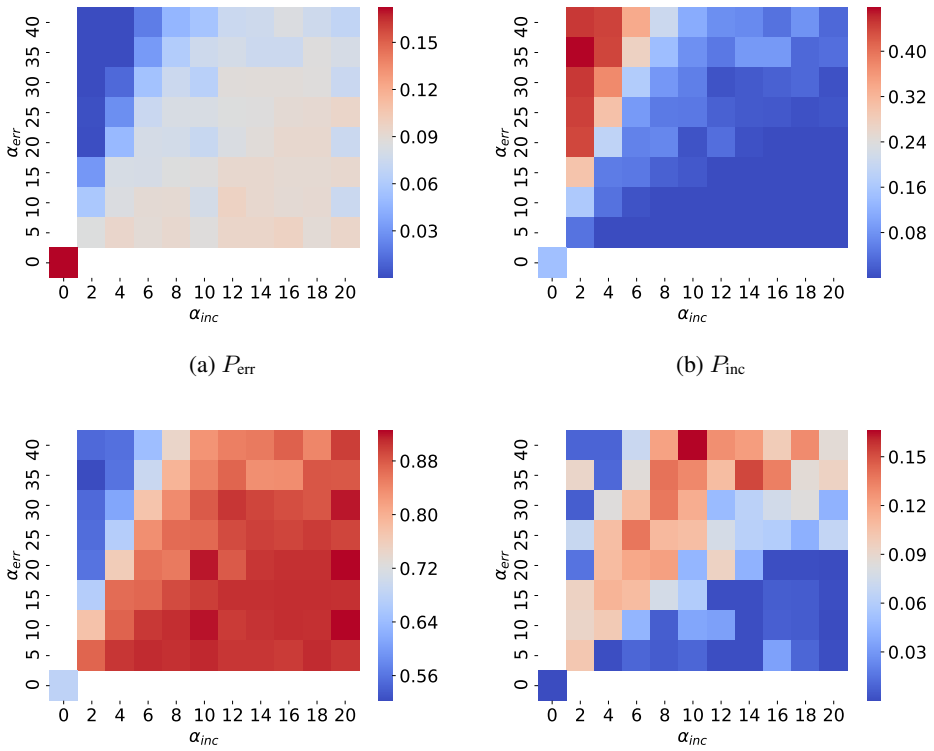


(a) $P_{\text{err}}$ (b) $P_{\text{inc}}$

(c) $P_{\text{suc}}$ (d) $P_{\text{suc}}$ SD

Figure 3: **With different penalties we observed a trade-off between the error rate and the inconclusive rate.** Compared with the point $\alpha_{\text{err}} = \alpha_{\text{inc}} = 0$ (bottom left corner), the added penalties improved the success probability or the inconclusiveness respectively. **(a)-(c)**: The gradual transition from the unambiguous classification (near-zero error rate, top left corner) to a minimal error classification (near-zero inconclusiveness, bottom right corner) with changes in the error penalty $\alpha_{\text{err}}$ and the inconclusiveness penalty $\alpha_{\text{inc}}$. We observed that the gain in the success rate was around $0.32$ with a sacrifice of only $0.1$ in the error rate. The data was tested on $a \in [0, 1]$, and averaged over $50$ repeated trails with random initializations. **(d)**: Standard deviation for $P_{\text{suc}}$. With an increasing standard deviation (closer to the diagonal line), the result became increasingly unstable when the two penalties ($\alpha_{\text{err}}$ and $\alpha_{\text{inc}}$) are closer in value. The standard deviations for $P_{\text{err}}$ and $P_{\text{inc}}$ showed the same pattern as for $P_{\text{suc}}$ (not shown).

Secondly, for all the dataset in Table 1, we achieved a much higher success rate than the theoretical case of an exactly $0$ error rate, when only requiring an error rate target being smaller than $0.01$ (Figure 4).

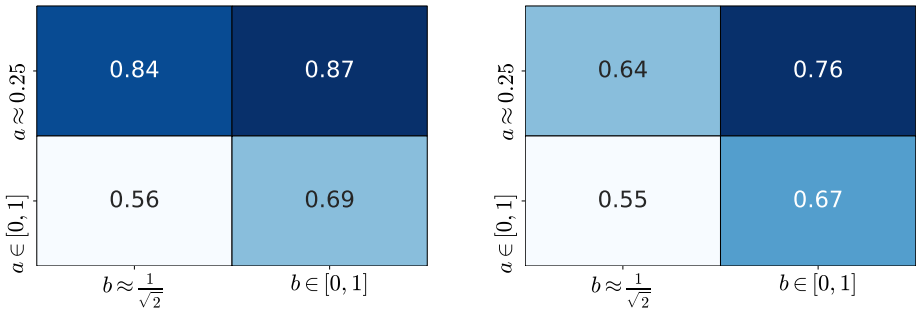


(a) Success rate for trained circuit for classification of data

(b) The theoretical success rate when exact 0 error rate is achieved.

Figure 4: **Unambiguous discrimination of data sampled from different non-orthogonal quantum probability distributions with higher success rate. (a)** Trained quantum circuits were capable of generalizing on quantum data which was sampled from a variety of different mixed probability distributions for multiple inputs. The classification was done in an unambiguous manner (with error rate $< 0.01$). **(b)** For comparison, we included here the theoretical results mentioned in Table 2.

## 5 Learning convergence from ensemble measurements

Although we trained our circuit with probabilities of various outcomes, which are not available on an actual quantum device (but we can assess these during the simulation), we found that using probabilities which are inferred from repeated measurements of the quantum states, the training process still convergences in all cases. In practice, hte noise introduced by the inferred probabilities could be effectively countered by increading the number of repeated measurements, using a lowered error rate, and adjusting the step size for calculation of gradients using the forward difference formula. Overall, a combination of a $0.01$ step size, a $10^5$ repeated measurements, and a learning rate of $10^{-3}$ well approximated the result obtained using exact probabilities and therefore our study here is feasible for actual quatum devices and we leave ofpen the actual implementation as future project. Further discussions are deferred in Appendix B.

## 6 Conclusions

We have developed an universal quantum circuit learning approach for classification of quantum data. In particular, we have designed a theoretically motivated loss function and used the stochastic optimization algorithm Adam in a quantum-classical hybrid scheme to train a circuit to perform quantum state discrimination. This training process generalized well for the discrimination task on new data, i.e., states from the parameter range which have not been seen during the training process. This in particular distinguishes our work from previous results on quantum circuit learning, in particular very recent study in e.g. Fanizza et al. (2018), which only optimizes circuits for specific inputs. Note that this prior work hence does not consider the generalization ability and hence does not treat the actual learning problem which is aiming at optimization as well as generalization.
We observed a trade-off between error and inconclusiveness rates when we penalized them differently in the loss function. Although this experiment was done on a simulated quantum computers, i.e., classical hardware, where exact measurement probabilities are available, we showed that this optimization could be experimentally performed with a repeated number of measurements of the wavefunction. Finally we note that recent quantum methods for estimating the analytical gradient via variations in the unitaries (Mitarai et al. (2018)) can be directly applied for training our circuits and therefore one can perform the optimization efficiently on near term quantum devices.

The discriminative quantum neural networks of the forms that were trained here could be potentially used as quantum shallow circuits for verifying or certifying other shallow or deep quantum circuits within machine learning or quantum simulations applications. They could also be used to verify the output of other generative models, such as Restricted Boltzmann machines or GANS Goodfellow

et al. (2014). Small-scale discriminative quantum circuit learning could be used for constructing non-trivial (e.g., POVM-based) receivers in quantum meteorology Giovannetti et al. (2011), sensing Degen et al. (2017), and imaging Chen et al. (2012).

## 7 ACKNOWLEDGEMENTS

We want to thank Raban Iten, Oliver Reardon-Smith and Roger Colbeck for valuable insights in parametrizing the general measurement circuits and Jarrod McClean for feedback on the manuscript. This project acknowledges the use of the ESPRC funded Tier 2 facility JADE, the use of the UCL Legion High Performance Computing Facility (Legion@UCL), and associated support services, in the completion. L.W. is supported by the Royal Society. S.S. is supported by the Royal Society, EPSRC, the National Natural Science Foundation of China, and the grant ARO-MURI W911NF-17-1-0304 (US DOD, UK MOD and UK EPSRC under the Multidisciplinary University Research Initiative). This work has been carried out while L.W. and S.S. participated in the workshop of *Measurement and control of quantum systems* at the Institut Henri Poincare. The financial support is kindly acknowledged.

## REFERENCES


Scott Aaronson. Read the fine print. *Nature Physics*, 11(4):291–293, 2015.

Andris Ambainis. Variable time amplitude amplification and a faster quantum algorithm for solving systems of linear equations. *arXiv preprint arXiv:1010.4458*, 2010.

Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *International Conference on Machine Learning*, pp. 1120–1128, 2016.

Leonardo Banchi, Nicola Pancotti, and Sougato Bose. Quantum gate learning in qubit networks: Toffoli gate without time-dependent control. *Npj Quantum Information*, 2:16019, 2016.

Stephen M. Barnett and Sarah Croke. Quantum state discrimination. *Advances in Optics and Photonics*, 1(2):238, apr 2009. ISSN 1943-8206. doi: 10.1364/AOP.1.000238. URL https://www.osapublishing.org/aop/abstract.cfm?uri=aop-1-2-238.

Charles H Bennett. Quantum cryptography using any two nonorthogonal states. *Physical review letters*, 68(21):3121, 1992.

János A Bergou. Quantum state discrimination and selected applications. *Journal of Physics: Conference Series*, 84:012001, oct 2007. ISSN 1742-6596. doi: 10.1088/1742-6596/84/1/012001. URL http://stacks.iop.org/1742-6596/84/i=1/a=012001?key=crossref.675d05b3d19f4f2bf39356d3f466d6b1.

Jacob Biamonte, Peter Wittek, Nicola Pancotti, Patrick Rebentrost, Nathan Wiebe, and Seth Lloyd. Quantum machine learning. *Nature*, 549(7671):195, 2017.

Anthony Chefles. Quantum state discrimination. *Contemporary Physics*, 41(6):401–424, 2000.

Jian Chen, Jonathan L. Habif, Zachary Dutton, Richard Lazarus, and Saikat Guha. Optical codeword demodulation with error rates below the standard quantum limit using a conditional nulling receiver. *Nature Photonics*, 6:374 EP –, 05 2012. URL http://dx.doi.org/10.1038/nphoton.2012.113.

Andrew M Childs. Quantum algorithms: Equation solving by simulation. *Nature Physics*, 5(12): 861–861, 2009.

Andrew M Childs, Robin Kothari, and Rolando D Somma. Quantum algorithm for systems of linear equations with exponentially improved dependence on precision. *SIAM Journal on Computing*, 46(6):1920–1950, 2017.

Carlo Ciliberto, Mark Herbster, Alessandro Davide Ialongo, Massimiliano Pontil, Andrea Rocchetto, Simone Severini, and Leonard Wossnig. Quantum machine learning: a classical perspective. *Proc. R. Soc. A*, 474(2209):20170551, 2018.

C. L. Degen, F. Reinhard, and P. Cappellaro. Quantum sensing. *Rev. Mod. Phys.*, 89:035002, Jul 2017. doi: 10.1103/RevModPhys.89.035002. URL https://link.aps.org/doi/10.1103/RevModPhys.89.035002.

Danial Dervovic, Mark Herbster, Peter Mountney, Simone Severini, Naïri Usher, and Leonard Wossnig. Quantum linear systems algorithms: a primer. *arXiv preprint arXiv:1802.08227*, 2018.

Lu-Ming Duan and Guang-Can Guo. Probabilistic cloning and identification of linearly independent quantum states. *Physical review letters*, 80(22):4999, 1998.

Marco Fanizza, Andrea Mari, and Vittorio Giovannetti. Optimal universal learning machines for quantum state discrimination. *arXiv preprint arXiv:1805.03477*, 2018.

Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.

Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Advances in quantum metrology. *Nature Photonics*, 5:222 EP –, 03 2011. URL http://dx.doi.org/10.1038/nphoton.2011.35.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pp. 2672–2680, 2014.

Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G Green, and Simone Severini. Hierarchical quantum classifiers. *arXiv preprint arXiv:1804.03680*, 2018.

Aram W Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical review letters*, 103(15):150502, 2009.

Stephanie L Hyland and Gunnar Rätsch. Learning unitary operators with help from u (n). In *AAAI*, pp. 2050–2058, 2017.

Luca Innocenti, Leonardo Banchi, Alessandro Ferraro, Sougato Bose, and Mauro Paternostro. Supervised learning of time-independent hamiltonians for gate design. *arXiv preprint arXiv:1803.07119*, 2018.

Raban Iten, Roger Colbeck, Ivan Kukuljan, Jonathan Home, and Matthias Christandl. Quantum Circuits for Isometries. *Physical Review A - Atomic, Molecular, and Optical Physics*, jan 2015. doi: 10.1103/PhysRevA.93.032318. URL http://arxiv.org/abs/1501.06911http://dx.doi.org/10.1103/PhysRevA.93.032318.

Raban Iten, Roger Colbeck, and Matthias Christandl. Quantum Circuits for Quantum Channels. *Physical Review A - Atomic, Molecular, and Optical Physics*, 93 (3):052316, sep 2016. ISSN 24699934. doi: 10.1103/PhysRevA.95.052316. URL http://arxiv.org/abs/1609.08103{%}0Ahttp://link.aps.org/doi/10.1103/PhysRevA.95.052316http://arxiv.org/abs/1609.08103http://dx.doi.org/10.1103/PhysRevA.95.052316.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. dec 2014. URL http://arxiv.org/abs/1412.6980.

Ying Li and Simon C Benjamin. Efficient variational quantum simulator incorporating active error minimization. *Physical Review X*, 7(2):021050, 2017.

Seth Lloyd. Enhanced sensitivity of photodetection via quantum illumination. *Science*, 321(5895): 1463–1465, 2008. ISSN 0036-8075. doi: 10.1126/science.1160627. URL http://science.sciencemag.org/content/321/5895/1463.

Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411*, 2013.

Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. Quantum principal component analysis. *Nature Physics*, 10(9):631–633, 2014.

Kosuke Mitarai, Makoto Negoro, Masahiro Kitagawa, and Keisuke Fujii. Quantum circuit learning. *arXiv preprint arXiv:1803.00745*, 2018.

Masoud Mohseni, Aephraim M. Steinberg, and János A. Bergou. Optical Realization of Optimal Unambiguous Discrimination for Pure and Mixed Quantum States. *Physical Review Letters*, 93 (20):200403, nov 2004. ISSN 0031-9007. doi: 10.1103/PhysRevLett.93.200403. URL https://link.aps.org/doi/10.1103/PhysRevLett.93.200403.

Philippe Raynal, Norbert Lütkenhaus, and Steven J. van Enk. Reduction theorems for optimal unambiguous state discrimination of density matrices. *Phys. Rev. A*, 68:022308, Aug 2003. doi: 10.1103/PhysRevA.68.022308. URL https://link.aps.org/doi/10.1103/PhysRevA.68.022308.

Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. Quantum support vector machine for big data classification. *Physical review letters*, 113(13):130503, 2014.

Andrea Rocchetto, Edward Grant, Sergii Strelchuk, Giuseppe Carleo, and Simone Severini. Learning hard quantum distributions with variational autoencoders. 2017. doi: 10.1038/s41534-018-0077-z.

Jonathan Romero, Jonathan P Olson, and Alan Aspuru-Guzik. Quantum autoencoders for efficient compression of quantum data. *Quantum Science and Technology*, 2(4):045001, 2017.

Andrew M Saxe, James L McClelland, and Surya Ganguli. Exact solutions to the nonlinear dynamics of learning in deep linear neural networks. *arXiv preprint arXiv:1312.6120*, 2013.

Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. An introduction to quantum machine learning. *Contemporary Physics*, 56(2):172–185, 2015.

Maria Schuld, Alex Bocharov, Krysta Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *arXiv preprint arXiv:1804.00633*, 2018.

Vivek V. Shende, Igor L. Markov, and Stephen S. Bullock. Smaller two-qubit circuits for quantum communication and computation. *Proceedings - Design, Automation and Test in Europe Conference and Exhibition*, 2:980–985, 2004. ISSN 1530-1591. doi: 10.1109/DATE.2004.1269020.

Vivek V. Shende, Stephen S. Bullock, and Igor L. Markov. Synthesis of Quantum Logic Circuits. *Ieee Transactions on Computer-Aided Design of Integrated Circuits and Systems*, pp. 18, 2006. ISSN 0278-0070. doi: 10.1109/TCAD.2005.855930. URL http://arxiv.org/abs/quant-ph/0406176.

Casper Kaae Snderby, Tapani Raiko, Lars Maale, Sren Kaae Snderby, and Ole Winther. Ladder variational autoencoders, 2016.

Guillaume Verdon, Michael Broughton, and Jacob Biamonte. A quantum algorithm to train neural networks using low-depth circuits. *arXiv preprint arXiv:1712.05304*, 2017.

Kwok Ho Wan, Oscar Dahlsten, Hlér Kristjánsson, Robert Gardner, and MS Kim. Quantum generalisation of feedforward neural networks. *npj Quantum Information*, 3(1):36, 2017.

Nathan Wiebe, Ashish Kapoor, Christopher Granade, and Krysta M Svore. Quantum inspired training for boltzmann machines. *arXiv preprint arXiv:1507.02642*, 2015.

Leonard Wossnig, Zhikuan Zhao, and Anupam Prakash. A quantum linear system algorithm for dense matrices. *arXiv preprint arXiv:1704.06174*, 2017.

# A QUANTUM CIRCUITS FOR POVM

This section describes the parametrization of the circuit capable of performing any quantum measurement on 2 qubit inputs with 4 possible measurement outcomes. This circuit could be represented by the following circuit diagram:

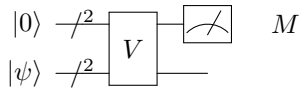


## A.1 COSINE-SINE DECOMPOSITION

Here we mention the Cosine-sine decomposition of unitary matrix, which will be frequently used in the following. For every unitary matrix $U \in \mathbb{C}^{2^n \times 2^n}$, it can be decomposed as:

$$U_n = \begin{pmatrix} A_0 & 0 \\ 0 & A_1 \end{pmatrix} \begin{pmatrix} C & -S \\ S & C \end{pmatrix} \begin{pmatrix} B_0 & 0 \\ 0 & B_1 \end{pmatrix} \tag{A.1}$$

where $A_0, A_1, B_0, B_1$ are unitary matrices of size $2^{n-1} \times 2^{n-1}$, $C$ and $S$ are real diagonal matrices of size $2^{n-1} \times 2^{n-1}$ satisfying $C^2 + S^2 = \mathbb{1}$. It is written in the following circuit equivalence:

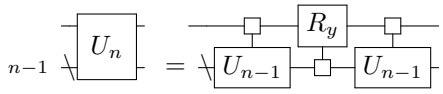


(A.2)

Where a box $\square$ represents the control part of a uniformly controlled gate, see section IV of Iten et al. (2015) for details. In the circuit in Fig. A, the first qubit is initiated to be $|0\rangle$, so we have

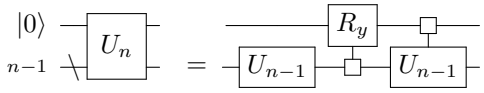


(A.3)

## A.2 DECOMPOSITION OF CIRCUIT IN FIG. A

For a general measurement giving at most 4 measurement outcomes, we have the following circuit representation:

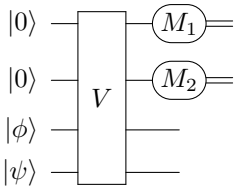


(A.4)

The first $V$ could be decomposed using the circuit equivalence on page 5 of Iten et al. (2016) into:

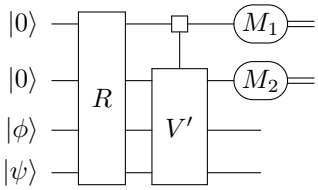


where the $R$ gate does not act on the second qubit. Applying Cosine-sine decomposition gives

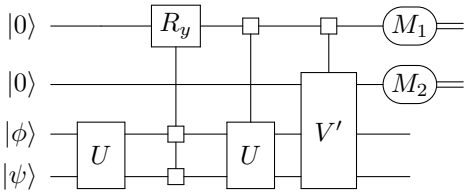

The uniform controlled $V'$ and $U$ can be merged and put after measurement of $M_1$ as:

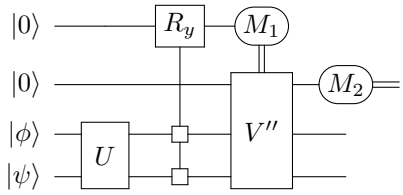


The first line of the circuit could be merged with the second line as follows:

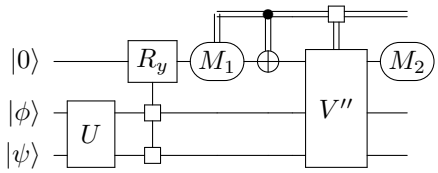


(A.5)

And then we apply the Cosine-sine decomposition to $V''$, throwing away the last gate on third and fourth qubits, we obtain:

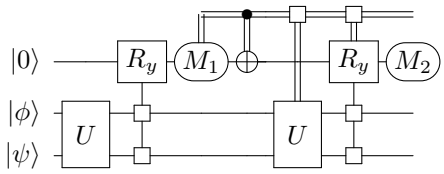


(A.6)

The uniformly-controlled rotations and remaining two qubit unitary gates could be easily parametrised by CNOTs and single qubit rotations. For example, see Shende et al. (2006) and Shende et al. (2004).

## B Learning convergence from ensemble measurements

Here we simulate the process that a classical-quantum hybrid scheme would implement utilizing a quantum device and analyse its performance. These numerical simulations can in principle be validated in a physical experiment, where the measurement outcomes are used to infer the different probabilities for the cost function. To have good estimation of the probability, and hence the cost function, one has to make repeated measurements to train the model, and we note that in particular better methods to evaluate the analytical gradient are available on a shallow quantum device Mitarai et al. (2018). We first give a brief discussion of the an estimated number of repeated measurements which are required to approximate the gradient, which is oriented on Farhi & Neven (2018)[Section 3]. Since the gradients are calculated using the forward difference formula:

$$\frac{df}{dx}(x) = \frac{f(x+\varepsilon) - f(x)}{\varepsilon} + O(\varepsilon^2) \tag{B.1}$$

The error in the calculation of $f$ must be at most of the order of $O(\varepsilon^2)$, in order to prevent dominating the total error. To achieve this ideally with an 99% probability, one requires the number of repeated measurement to be of the order $\frac{1}{(\varepsilon^2)^2} = \frac{1}{\varepsilon^4}$.[1] For example, when $\varepsilon = 10^{-3}$, the ideal number of repetitions is given by $10^{12}$.

In practice, we do not use $\frac{1}{\varepsilon^4}$ measurements, since Adam is designed with the stochasticity of cost function taken into account. To give an estimate of the number of repeated measurements which are required for convergence of the optimisation process, we perform two numerical experiments. We first look at the case when the number of repeated measurements is $\geq 10^3$ and $\varepsilon = 10^{-2}$.

[1] This assumes that the cost function follows a normal distribution with variance of the order $\frac{1}{\sqrt{N}}$, where $N$ is the number of measurements made in reach run in order to calculate the cost function.

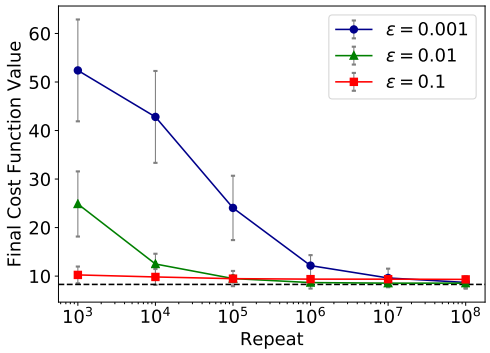


Figure 5: The cost function after 5000 iterations. The result obtained using exact probabilities is shown by the horizontal dashed line. For a smaller step size for gradient ($\varepsilon$), we find that more repetitions are required to give a consistent result. However, a combination of $\varepsilon = 10^{-2}$ and $10^5$ repetitions gives a result which well approximates the result obtained using exact probabilities. Here *repeat* is the number of repeated measurements that are made each time to calculate the cost function. The cost function values are averaged over 50 repeated runs of the training process, and the bars indicate the standard deviations.

We find that $10^5$ repeated measurements for each iteration are a robust configuration for successful convergence. Second, we vary the learning rate and increase the maximal number of iterations for Adam, setting $\varepsilon = 10^{-2}$ and taking only 100 repeated measurements. We observe that optimisation is successful with the large number of iteration. In both experiments, the penalties were set to $\alpha_{inc} = 5$ and $\alpha_{err} = 40$.

**Large number of repetitions.** Our results show that for a fixed maximum number of iterations (5000) for Adam, a combination of $\varepsilon = 10^{-2}$ and $10^5$ repeated measurements gives robust results, i.e., the final cost function is close to the value obtained with the exact probability (with an error within $3\%$) and is stable (with a relative standard deviation of $13\%$). A more detailed description of the trade-off between repeated measurements and the stability of the cost function is shown in Fig. 5.

**Small learning rates and high number of iterations.** Our numerical experiments further showed that in the case of small repetitions, lowered learning rate could effectively counter the noisy brought by the insufficient sampling. Although in this case, the optimisation required a longer iteration to finish. For example, with only 100 repeated measurements, the variance of cost function $J_1$ after 20000 iterations decreased as we lowered the learning rate (Fig.6(a)). We could visually observe the optimisation process where the cost function $J_1$ slowly approached the optimal value in Fig 6(b). Here, gradient step were taken as $\varepsilon = 10^{-2}$.

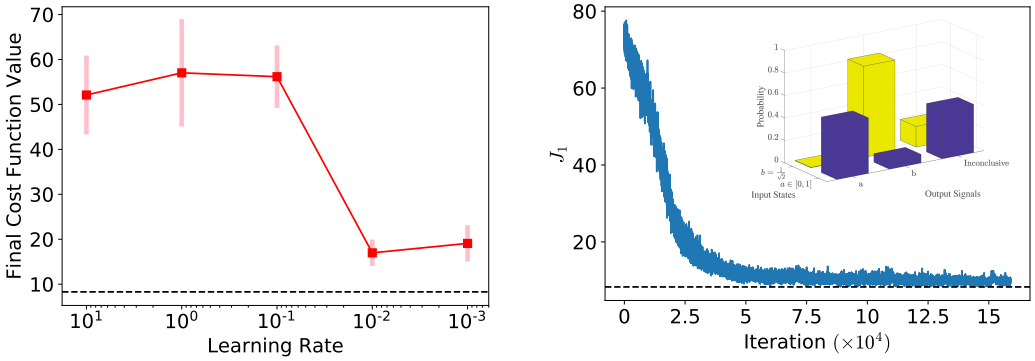


(a) Lowering learning rate to counter the effect of insuffi-cient sampling. Both the value and the standard deivationmeasurements slowly moved to its optimal value when of cost function at the 20000'th iteration were brounghtthe learning rate is set as $0.001$. The horizontal dashed down by lowered learning rate. Here the number of re-line showed the minimal value $8.3$ for cost function. The peated measurements is only $100$, much smaller thaninset illustrate that trained circuit could discrminate the $\frac{1}{\varepsilon^4} = 10^8$. two quantum states. Although the error rate in the inset is not $0$, we believe it could be achieved by further tuning the penalities, which was set as ($\alpha_{err} = 40$, $\alpha_{inc} = 5$) for this training.

Figure 6: Small learning rates with high number of iterations.