

VSQ: Variational Shadow Quantum Learning for Classification

Guangxi Li,^{1,2} Zhixin Song,¹ Xin Wang^{1*}

¹Institute for Quantum Computing, Baidu Research, Beijing 100193, China

²Centre for Quantum Software and Information, University of Technology Sydney, NSW 2007, Australia
 guangxi.li@student.uts.edu.au, zhixinsong0524@gmail.com, wangxin73@baidu.com

Abstract

Classification of quantum data is essential for quantum machine learning and near-term quantum technologies. In this paper, we propose a new hybrid quantum-classical framework for supervised quantum learning, which we call Variational Shadow Quantum Learning (VSQ). Our method in particular utilizes the classical shadows of quantum data, which fundamentally represent the side information of quantum data with respect to certain physical observables. Specifically, we first use variational shadow quantum circuits to extract classical features in a convolution way and then utilize a fully-connected neural network to complete the classification task. We show that this method could sharply reduce the number of parameters and thus better facilitate quantum circuit training. Simultaneously, less noise will be introduced since fewer quantum gates are employed in such shadow circuits. Moreover, we show that the Barren Plateau issue, a significant gradient vanishing problem in quantum machine learning, could be avoided in VSQ. Finally, we demonstrate the efficiency of VSQ in quantum classification via numerical experiments on the classification of quantum states and the recognition of multi-labeled handwritten digits. In particular, our VSQ approach outperforms existing variational quantum classifiers in the test accuracy in the binary case of handwritten digit recognition and notably requires much fewer parameters.

Introduction

Quantum computers are expected to have significant applications in solving challenging problems in information processing. Inspired by the powerful capacity of classical supervised learning and its growing community (Lecun, Bengio, and Hinton 2015; Goodfellow, Bengio, and Courville 2016), it is natural to develop their quantum counterparts and explore the emerging field of quantum machine learning (QML) (Biamonte et al. 2017; Schuld, Sinayskiy, and Petruccione 2015; Ciliberto et al. 2018; Lloyd, Mohseni, and Rebentrost 2013). Among many topics in this area, classification is one of the most important tasks, e.g., distinguishing quantum states (Bae and Kwek 2015; Chen et al. 2018; Patterson et al. 2019) or recognizing classical data (Havlíček et al. 2019; Benedetti et al. 2019; Schuld et al. 2020). Classification is usually described as a decision making process

with discrete variable where the processing unit is provided with a labeled training set $\mathcal{D}^{(train)} = \{(\rho^{(m)}, y^{(m)})\}$ in order to find the convoluted mapping \mathcal{F} between each set element $\rho^{(m)}$ and its corresponding label $y^{(m)}$. Once the training process is complete, we would expect the classifier \mathcal{F} not only learns the map $\mathcal{F}(\rho^{(m)}) = y^{(m)}$ precisely, but also generalizes its capacity of discrimination to discover some hidden features shared with similar test data $\mathcal{F}(\rho^{(new)}) = y^{(new)}$ (i.e. recognize an unseen cat as a cat). This ability of generalization is valuable to all classification tasks, and hence it is frequently used to benchmark the performance of a classifier.

In classical machine learning, various approaches have been proposed to implement classification tasks, including perceptron-based algorithms, support vector machines, and the most prevalent neural network (NN) framework (Lecun, Bengio, and Hinton 2015). With the quantum computing community growing in the NISQ era (Preskill 2018), similar ideas have been developed respectively, including the quantum perceptron model (Kapoor, Wiebe, and Svore 2016), kernel-based method (Li, Chakrabarti, and Wu 2019), and the quantum neural network (QNN) framework (Havlíček et al. 2019; Mitarai et al. 2018; Schuld et al. 2020; Grant et al. 2018; Farhi and Neven 2018; Schuld and Killoran 2019; Schuld and Petruccione 2018; Bhatia et al. 2019).

This paper focuses on the QNN-based algorithms, also referred to as Variational Quantum Algorithms (VQA) or hybrid quantum-classical algorithms since they are regarded as well-suited for execution on NISQ devices by combining quantum computers with classical computers. The main idea of VQA is employing parametrized quantum circuits (as a unitary neural network architecture) to search the Hilbert space and combining classical optimization methods like gradient descent (GD) to find the best parameters (Peruzzo et al. 2014; Kandala et al. 2017; Farhi, Goldstone, and Gutmann 2014; Farhi and Harrow 2016). VQA has been applied to many topics such as quantum eigensolver (Peruzzo et al. 2014), quantum simulation (Yuan et al. 2019), quantum state distance estimation (Cerezo et al. 2020a; Chen et al. 2020) and quantum matrix decomposition (Wang, Song, and Wang 2020). So far, most proposals for variational quantum classification process information in the global sense such that the quantum circuit always acts on the whole Hilbert space fulfilling high-dimensional transformation. And the classi-

*Corresponding author.

cal feature/information extracted from the quantum system is achieved through measurement. This formulation faces two potential challenges. There could exist more efficient architectures to achieve the same task performance but significantly reducing the quantum resource required (number of quantum gates) by limiting the operating scope to a few selected qubits. The other challenge is the notorious Barren Plateau problem (McClean et al. 2018). As the problem size increases, it will exhibit exponentially vanishing gradients, making the optimization landscape flat and hence untrainable using gradient-based optimization methods.

To overcome the above challenges, we explore a significantly different hybrid architecture based on classical shadows. In comparison, our method extracts only "local" features from the subspace, which we call *shadow features*. We introduce shadow features for quantum classification due to the inspiration from the recent studies on shadows (Aaronson 2018; Huang, Kueng, and Preskill 2020). Specifically, important quantum properties such as quantum fidelities and entanglement entropies can be predicted using classical shadows rather than possessing full information of the quantum states (Huang, Kueng, and Preskill 2020). This provides us the intuition that classical shadows may also be applied in quantum classification.

Contributions. Our work is, as far as we know, the first attempt to combine the concept of shadow with the mindset of learning. We propose a Variational Shadow Quantum Learning (VSQL) framework that could be adapted to many near-term quantum applications. In particular, we apply this framework to develop quantum classifiers for near-term quantum devices. Firstly, we employ the parameterized shadow quantum circuits $U(\theta)$ (denoted as *shadow circuits*) acting on selected local qubit subspace rather than the whole qubit Hilbert space, which considers the operating scope efficiency and the connectivity limit on quantum hardware. Secondly, the shadow features of the input data (encoded as quantum states $\rho^{(m)}$ with labels $y^{(m)}$) will be computed via measuring the Pauli $X \otimes X \cdots \otimes X$ observables on the quantum devices. The final step is to utilize a classical Fully-Connected Neural Network (FCNN) to post-process these shadow features, and we could then decide the label prediction $\hat{y}^{(m)}$ through an activation function. We refer to Fig. 1 for a detailed sketch of our method for binary classification.

The contributions of our work are multi-folded. First, we introduce a hybrid quantum-classical framework that can be easily implemented on quantum devices with topological connectivity limits, since our VSQL mainly considers locally-operated quantum circuits. Second, we show that VSQL involves significantly fewer number of parameters (independent of the problem size) than existing variational quantum classifiers (Schuld et al. 2020; Farhi and Neven 2018). Notably, we prove that VSQL could naturally avoid the Barren Plateau issue (McClean et al. 2018) (gradients vanishing issue in QML) by limiting the operating scope. Finally, we demonstrate real-world applications of VSQL to do quantum state classification and handwritten digits recognition. We in particular show that VSQL outperforms existing variational quantum classifiers in the test accuracy while requiring much fewer parameters.

Method

Preliminaries and Notations

Here, we briefly introduce the basic concepts of quantum computation that are necessary for this paper. Interested readers are recommended to the celebrated textbook by Nielsen and Chuang (Nielsen and Chuang 2002). Information in the quantum computing field is represented by n -qubit quantum states over space $\mathbb{C}^{2^n \times 2^n}$, which could be mathematically described by positive semi-definite matrices $\rho \succeq 0$ with property $\text{Tr}(\rho) = 1$. Following this density matrix formulation, a quantum state is pure if $\text{Rank}(\rho) = 1$; otherwise, it is mixed. For a pure state ρ , it can be represented by a unit vector in the sense that $\rho = |\psi\rangle\langle\psi|$, where the *ket* notation $|\psi\rangle \in \mathbb{C}^d$ denotes a column vector and *bra* notation $\langle\psi| = |\psi\rangle^\dagger$ with \dagger denoting conjugate transpose. In general, we would also use $|\psi\rangle$ to denote a pure state for simplicity. A mixed state could be represented as $\rho = \sum_i q_i |\psi_i\rangle\langle\psi_i|$, where the coefficients $q_i \geq 0$ records the probability for quantum system to be in each corresponding pure state $|\psi_i\rangle\langle\psi_i|$ and hence $\sum_i q_i = 1$.

In this paper, the unitary matrices

$$X := \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, Y := \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, Z := \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (1)$$

refer specifically to the Pauli matrices. Their corresponding rotation gates are denoted by $R_P(\theta) := e^{-i\theta P/2} = \cos\frac{\theta}{2}\mathbb{I} - i\sin\frac{\theta}{2}P$, where $\theta \in [0, 2\pi]$ is the rotation angle and $P \in \{X, Y, Z\}$. The evolution of a quantum state ρ could be mathematically described by employing a quantum circuit (or QNN) $\rho' = U\rho U^\dagger$, where the unitary U is usually parameterized by a series of rotation gate angles θ and basic two-qubit gates. Quantum measurement is usually introduced at the end of algorithms to extract classical information from the quantum states. One hardware-friendly way to fulfill this is through the Pauli expectations $\langle P \rangle = \text{Tr}(P\rho')$, where $P \in \{X, Y, Z\}^{\otimes n}$ is a single Pauli matrix or Pauli product matrix, e.g., $P = X \otimes X \cdots \otimes X$.

Sketch of Our Method

We now present the sketch of VSQL for binary classification. Our goal is to find the optimal parameters θ^* in the shadow circuits $U(\theta)$ and the best weights $\{\mathbf{w}^*, b^*\}$ in the FCNN such that the algorithm could correctly predict the label of an unknown input quantum state. Like most classifiers, VSQL consists of two separate processes, training and inference. During the training process illustrated in Fig. 1, we are given the training data set encoded in n -qubit quantum state $\mathcal{D}^{(train)} := \{(\rho_{in}^{(m)}, y^{(m)})\}_{m=1}^{N_{train}}$, where $y^{(m)} \in \{0, 1\}$ denotes the binary label for the m -th input density matrix $\rho_{in}^{(m)}$. Then, the n_{qsc} -local shadow circuit acts on the first n_{qsc} qubits and the Pauli- $(X \otimes \cdots \otimes X)$ expectation value is estimated, recorded as *shadow feature* o_1 . Next, the identical shadow circuit is implemented on the subspace spanned from the 2^{nd} up to the $(2 + n_{qsc} - 1)^{th}$ qubit to extract the second shadow feature o_2 . As the shadow circuit sliding down, we obtain $n - n_{qsc} + 1$ shadow features in total. This convolution-like way of sliding through the

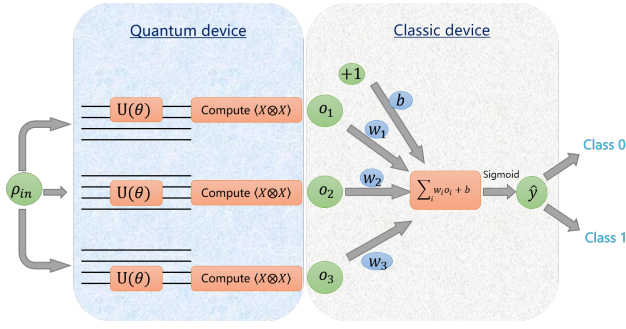


Figure 1: Sketch of variational shadow quantum learning (VSQL) for binary classification with $n = 4$ and $n_{qsc} = 2$. In the quantum device, the shadow circuit is implemented on the subspace of input state ρ_{in} . Sliding through the whole system to collect the Pauli- $(X \otimes X)$ expectations, i.e., shadow features. In the classic device, the resulting shadow features o_i 's are fed into a fully-connected neural network. Here, the output \hat{y} is a value between 0 and 1 for binary case. We should denote that all the shadow circuits $U(\theta)$'s sliding through the n -qubit Hilbert space are identical.

qubit positions can be adjusted according to the hardware connectivity. We also note that there is only one shadow circuit here. However, the number of shadow circuits (n_s) could be increased appropriately to accommodate the difficulty of classification tasks, with $n_s(n - n_{qsc} + 1)$ shadow features. Sequentially, we feed these local features $\{o_i\}$ into a classical FCNN, which means they are weighted summed with weights $\mathbf{w} \in \mathbb{R}^{n - n_{qsc} + 1}$, bias $b \in \mathbb{R}$ and mapped into the range $\hat{y}^{(m)} \in [0, 1]$ via the sigmoid activation function $\sigma(z) = (1 + e^{-z})^{-1}$. Repeat the same procedure for each input data and compute the accumulated loss $\mathcal{L}(\boldsymbol{\theta}, \mathbf{w}, b; \mathcal{D}^{(train)})$ between the predicted value $\hat{y}^{(m)}$ and its true label $y^{(m)}$. Finally, VSQl utilizes gradient-based optimizer to update the shadow circuit parameters $\boldsymbol{\theta}$ and the neural network parameters \mathbf{w}, b and gradually minimizing the loss function. Repeat these steps until the loss is converged with tolerance $\Delta\mathcal{L} \leq \varepsilon$ or other stopping criteria is satisfied. Check Algorithm 1 for details.

During the inference process, the unseen test data set $\mathcal{D}^{(test)} := \{(\rho_{in}^{(m)}, y^{(m)}) \in \{0, 1\}\}_{m=1}^{N_{test}}$ is provided to the classifier \mathcal{F} . We feed each sample in the test set to the trained hybrid framework (combination of shadow circuit and FCNN) to predict its label. Then, the test accuracy could be calculated by comparing the predicted labels and the true labels. Due to the space limitation, we refer the details to Appendix. Furthermore, VSQl can be naturally generalized to multi-label classification by replacing the sigmoid activation function with a softmax function. We defer the details for multi-label classification to Appendix.

Loss Function

Given the data set $\mathcal{D} := \{(\rho_{in}^{(m)}, y^{(m)})\}_{m=1}^N$ and n_{qsc} -local shadow circuits, the loss function of VSQl for binary clas-

Algorithm 1 Variational shadow quantum learning (VSQl) for binary classification: the training process

- Input:** The training data set $\mathcal{D}^{(train)} := \{(\rho_{in}^{(m)}, y^{(m)}) \in \{0, 1\}\}_{m=1}^{N_{train}}$, $EPOCH$, optimization procedure
Output: The final parameters $\boldsymbol{\theta}^*$, \mathbf{w}^* and b^* , and the list of losses
- 1: Initialize the parameters $\boldsymbol{\theta}$ of the 2-local (for example) shadow circuit $U(\boldsymbol{\theta})$ from uniform distribution $\text{Uni}[0, 2\pi]$ and \mathbf{w}, b from Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbb{I})$
 - 2: **for** $ep = 1, \dots, EPOCH$ **do**
 - 3: **for** $m = 1, \dots, N_{train}$ **do**
 - 4: Apply multi-times the shadow circuit $U(\boldsymbol{\theta})$ to the input density matrix $\rho_{in}^{(m)}$
 - 5: Measure the subsystem and estimate a series of expectations $\langle X \otimes X \rangle$, recorded as o_i 's
 - 6: Feed the shadow features o_i 's into the classical neural network and obtain the output $\hat{y}^{(m)}$
 - 7: Compute the accumulated loss $(\hat{y}^{(m)} - y^{(m)})^2$ and update accordingly the parameters $\boldsymbol{\theta}$, \mathbf{w} and b via gradient-based optimization procedure
 - 8: **end for**
 - 9: **if** the stopping criterion is satisfied **then**
 - 10: Break
 - 11: **end if**
 - 12: **end for**

sification is designed to be the mean square error¹ (Ziegel, Lehmann, and Casella 1999):

$$\mathcal{L}(\boldsymbol{\theta}, \mathbf{w}, b; \mathcal{D}) := \frac{1}{2N} \sum_{m=1}^N [\hat{y}^{(m)}(\rho_{in}^{(m)}; \boldsymbol{\theta}, \mathbf{w}, b) - y^{(m)}]^2. \quad (2)$$

Here, the predicted label $\hat{y}^{(m)}$ is defined as follows:

$$\hat{y}^{(m)}(\rho_{in}^{(m)}; \boldsymbol{\theta}, \mathbf{w}, b) := \sigma\left(\sum_i w_i o_i^{(m)}(\rho_{in}^{(m)}; \boldsymbol{\theta}) + b\right), \quad (3)$$

where $\sigma(z)$ denotes the sigmoid activation function and the shadow features o_i are calculated through

$$o_i^{(m)}(\rho_{in}^{(m)}; \boldsymbol{\theta}) = \text{Tr}\left(\rho_{in}^{(m)}(\mathbb{I} \otimes \dots \otimes U^\dagger(\boldsymbol{\theta}) O U(\boldsymbol{\theta}) \otimes \dots \otimes \mathbb{I})\right), \quad (4)$$

Note that the shadow circuit $U(\boldsymbol{\theta})$ and the physical observable $O = X \otimes \dots \otimes X$ are applied on the same local qubits. Additionally, $U(\boldsymbol{\theta})$ is usually decomposed as a chain of unitary operators:

$$U(\boldsymbol{\theta}) = \prod_{l=L}^1 U_l(\theta_l) V_l. \quad (5)$$

where $U_l(\theta_l) = \exp(-i\theta_l P_l/2)$ and V_l denotes a fixed operator such as Identity, CNOT and so on.

¹The cross-entropy loss is considered in the multi-label case; see Appendix.

Analytical Gradients

With the above preparation, we can easily derive the analytical gradients, with which VSQL could naturally update its parameters θ and $\{w, b\}$ via gradient-based optimization method, e.g., SGD (Bottou 2004). For each input $\rho_{in}^{(m)}$,

$$\frac{\partial \mathcal{L}}{\partial w_i} = \left(\hat{y}^{(m)} - y^{(m)} \right) \cdot \hat{y}^{(m)} \left(1 - \hat{y}^{(m)} \right) \cdot o_i^{(m)}, \quad (6)$$

$$\frac{\partial \mathcal{L}}{\partial b} = \left(\hat{y}^{(m)} - y^{(m)} \right) \cdot \hat{y}^{(m)} \left(1 - \hat{y}^{(m)} \right), \quad (7)$$

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial \theta_l} &= \frac{\partial \mathcal{L}}{\partial \hat{y}^{(m)}} \cdot \sum_i \frac{\partial \hat{y}^{(m)}}{\partial o_i^{(m)}} \cdot \frac{\partial o_i^{(m)}}{\partial \theta_l} \\ &= \left(\hat{y}^{(m)} - y^{(m)} \right) \cdot \sum_i \hat{y}^{(m)} \left(1 - \hat{y}^{(m)} \right) w_i \cdot \frac{\partial o_i^{(m)}}{\partial \theta_l}, \quad (8) \end{aligned}$$

For the partial derivatives w.r.t w_i and b written in Eqs. (6) and (7), they would be directly computed in the classical device and used to update w_i, b through the back propagation algorithm (Goodfellow, Bengio, and Courville 2016). For the partial derivative w.r.t θ_l in Eq. (8), it can be regarded as a weighted sum of several partial gradients $\partial o_i^{(m)} / \partial \theta_l$,

$$\frac{\partial o_i^{(m)} \left(\theta; \rho_{in}^{(m)} \right)}{\partial \theta_l} = -\frac{i}{2} \text{Tr} \left(U_{>l}^\dagger O U_{>l} \left[P_l, U_{\leq l} \rho_i U_{\leq l}^\dagger \right] \right). \quad (9)$$

where $\rho_i = \text{Tr}_{-i}(\rho_{in}^{(m)})$ denotes the partial trace of $\rho_{in}^{(m)}$ corresponding to the index i , $U_{\leq l} = \prod_{j=1}^l U_j(\theta_j) V_j$ and $U_{>l} = \prod_{j=l+1}^L U_j(\theta_j) V_j$ and $[\rho, \sigma] = \rho\sigma - \sigma\rho$ is a commutator. This gradient can be calculated exactly on the quantum device with the $\pi/2$ parameter shift rule proposed by Mitarai et al.. Compared with the finite difference scheme, this method leads to a faster convergence (Harrow and Napp 2019) and is more suitable to the existing quantum devices.

Theoretical Performance Analysis

Number of Parameters of VSQL

In the hybrid quantum-classical framework, the number of parameters in the quantum circuit is an important metric to measure its complexity and efficiency. The main reason is that updating each parameter is costly in terms of quantum resources as it requires re-running the entire circuit multiple times. Therefore, algorithms with a smaller number of parameters are preferable in the NISQ era. Here, we will exhibit this advantage for VSQL.

There are two kinds of parameters in VSQL, i.e., the parameters θ in the shadow circuits and the parameters w, b in the classical NN. Assume the action mode of the shadow circuits is ‘‘shadow sliding’’ (illustrated in Fig. 1), which is also employed throughout this paper. The number of parameters of VSQL for binary classification is summarized as follows.

Proposition 1. *For an n -qubit quantum system, if we use n_s shadow circuits, then the number of parameters of VSQL for binary classification is*

$$\begin{aligned} \# \text{Params} &= \# \text{Params} \Big|_{\text{in shadow circuits}} + \# \text{Params} \Big|_{\text{in NN}} \\ &= n_s n_{qsc} D + [n_s (n - n_{qsc} + 1) + 1], \quad (10) \end{aligned}$$

where we denote by n_{qsc} the number of qubits of the shadow circuits and assume each shadow circuit consists of D layers with n_{qsc} parameters in each layer.

Thus, VSQL has a parameter quantity that is linearly related to n and D separately, rather than nD that commonly appears in most of the ansatzes employed in the existing literature (Schuld et al. 2020; Farhi and Neven 2018; Mitarai et al. 2018). For a 50-qubit quantum system, if we use just one 2-local shadow circuit with 20 layers, i.e., $n_s = 1, n_{qsc} = 2, D = 20$, then the number of parameters of VSQL is $40 + (50 - 2 + 1) + 1 = 90$, which is much smaller than $nD = 1000$.

Theoretical Classification Ability

In this subsection, we explore the theoretical classification ability of VSQL and give the corresponding necessary and sufficient conditions.

Theorem 2. *Given two types of input density matrices $\rho_{in}^{(0)}$ and $\rho_{in}^{(1)}$ with labels 0 and 1, respectively, if there exists a group of θ that makes at least one pair of shadow features $o_i^{(0)}$ and $o_i^{(1)}$ different, i.e., $|o_i^{(0)} - o_i^{(1)}| > 0$, then VSQL can distinguish them, vice versa.*

Corollary 2.1. *Given two types of n -qubit input density matrices. If each pair of their corresponding m -local partial traces are identical ($m < n$), then VSQL is theoretically incapable of distinguishing them via m -local shadow circuits, and vice versa.*

The proof of Corollary 2.1 could be immediately derived from Theorem 2, because getting identical partial traces is equivalent to having same shadow features (cf. Eq. (4)).

After exploring the necessary and sufficient condition for the theoretical classification ability of VSQL, we now discuss this ability under different local shadow circuits. Intuitively, larger shadow circuits will give VSQL stronger classification ability. The following Theorem will give a detailed statement.

Theorem 3. *Given two types of n -qubit input density matrices $\rho_{in}^{(0)}$ and $\rho_{in}^{(1)}$. If VSQL can not theoretically distinguish them via m -local shadow circuits, then neither can via m' -local shadow circuits, where $m' < m < n$. And not vice versa.*

The proof is shown in Appendix. From Theorem 3, we confirm the intuition that the larger the number of qubits n_{qsc} of the shadow circuits is, the stronger the theoretical expressive ability of VSQL is. However, if this number is too large, it will lead to other problems, such as the Barren Plateau issue described in the next subsection.

Escape of Barren Plateau

In the last subsection, we have shown that VSQL has a strong theoretical classification ability for a wide range of quantum states, especially by using large local shadow circuits. However, the sizeable operating scope of the shadow circuits will increase network parameters and the cost of compiling given limited hardware connections, and leads to the Barren Plateau (BP) issue. The BP issue (McClean et al.

2018; Cerezo et al. 2020b) refers to the vanishing gradient problem during the training process of QNN. That is, for a wide range of variational quantum circuits, the partial gradients of the objective function have a zero mean and an exponentially vanishing variance, which makes it difficult for the optimizer to find the correct direction to decrease the objective function. Therefore, it is important to discuss whether the BP problem exists when proposing a new variational quantum algorithm.

Next, we will evaluate the mean and the variance of the analytical gradients in VSQL. There is no BP issue for the partial gradients (see Eqs. (6) and (7)) with respect to the parameters w_i and b of the classical NN. And for the partial gradient (see Eq. (8)) with respect to θ_l of QNN, the BP issue is mainly reflected on the last term, i.e., the partial derivatives (see Eq. (9)) of the shadow features o_i with respect to θ_l . Hence, it is sufficient to evaluate the mean and the variance of the partial gradient in Eq. (9) to explore the BP problem in VSQL. The results are summarized in the following proposition.

Definition 4. A unitary t -design (Dankert et al. 2009) is defined as a finite set of unitaries $\{U_k\}_{k=1}^K$ on a d -dimensional Hilbert space such that

$$\frac{1}{K} \cdot \sum_k P_{(t,t)}(U_k) = \int_{\mathcal{U}(d)} d\mu_{\text{Haar}}(U) P_{(t,t)}(U), \quad (11)$$

where $P_{(t,t)}(U)$ denotes a polynomial of degree at most t on the elements of U and at most t on the elements of U^\dagger .

Proposition 5. If $U_{>1}$ or $U_{\leq 1}$ forms at least an n_{qsc} -local unitary 2-design, the mean and the variance of the analytical gradients with respect to θ_l in VSQL (see Eq. (9)) are evaluated as

$$\mathbb{E} \left[\frac{\partial o_i}{\partial \theta_l} \right] = 0; \quad \text{Var} \left[\frac{\partial o_i}{\partial \theta_l} \right] = -\frac{1}{4} \cdot \frac{C(\rho_i)}{2^{2n_{qsc}} - 1}, \quad (12)$$

where $C(\rho_i) \in (-4 \times 2^{n_{qsc}}, 0)$ denotes a constant and n_{qsc} is the number of qubits of the shadow circuits.

The proof is shown in Appendix. From Proposition 5, we notice the variance of the gradients decays exponentially with n_{qsc} , rather than the qubit number n . Hence, no matter how big the problem size n is, as long as we choose a small n_{qsc} (e.g., $n_{qsc} \leq 4$) and assume $C(\rho_i) \approx -2 \times 2^{n_{qsc}}$, we can evaluate the analytical gradients efficiently via more than 1,000 repetitions derived from the Chernoff bound. In one word, VSQL could escape the barren plateaus by choosing an appropriate operating scope n_{qsc} . Moreover, Wang et al. indicates that noise could also induce the BP issue. Following this line of reasoning, the small shadow circuits in VSQL, where less noise is introduced, will also be beneficial for escaping the barren plateaus from a different perspective.

Here, we provide an illustrated example. Assume the n -qubit quantum state $\rho_{in} = |\psi_{in}\rangle\langle\psi_{in}|$ we want to classify is labeled with 0, where

$$|\psi_{in}\rangle := \otimes_{j=0}^{n-1} R_y(2\pi j/n) |0\rangle. \quad (13)$$

The chosen shadow circuit consists of a layer of single-qubit R_y rotations and a layer of CNOT gates which only connects the adjacent qubits, followed by another layer of R_y

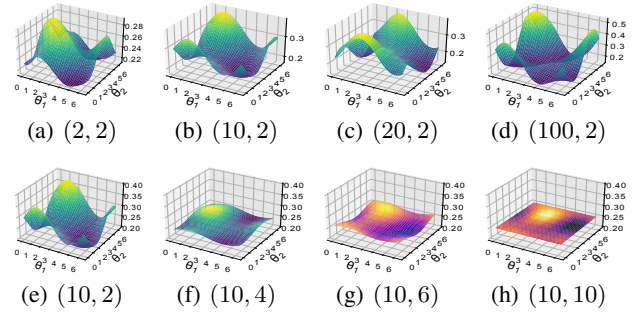


Figure 2: The slice of loss landscape with respect to the first two circuit parameters by changing the system size n and operating scope n_{qsc} . Here, the binary list represents (n, n_{qsc}) .

rotations. Then, we compute the loss landscape in Eq. (2) with regard to the first two circuit parameters by fixing all the other parameters with $\pi/4$ and setting the bias $b = 0$ and $\mathbf{w} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$ sampled from a Gaussian distribution. The result, as shown in Fig. 2, is in line with the above analysis, i.e., there is no barren plateaus with $n_{qsc} = 2$, but the loss landscape shrinks dramatically with an increasing n_{qsc} .

Numerical Experiments

We supplement our theoretical results with numerical experiments by classical simulation of VSQL. Specifically, our numerical experiments include distinguishing two (and three) families of 2-qubit quantum states and classifying handwritten digit images taken from the MNIST data set. We also conduct experiments on classifying noisy quantum states to exhibit the robustness of VSQL, which is deferred to Appendix due to limited space. All the simulations and optimization loop are implemented via Paddle Quantum² on the PaddlePaddle Deep Learning Platform (Ma et al. 2019).

Classification of Quantum States

Quantum state discrimination (QSD) is a fundamental problem in quantum information and it underlies various applications (Nielsen and Chuang 2002; Roa, Retamal, and Saavedra 2002; Bae and Kwek 2015). The goal of QSD is to determine which state a given physical system has been prepared. In general, finite copies of the given states are considered in the study of distinguishing quantum states (Lloyd et al. 2020; Gambs 2008).

Classification of Binary Quantum States. We choose two canonical families of non-orthogonal 2-qubit quantum states as a proof of principle. These states are well-studied in Refs. (Mohseni, Steinberg, and Bergou 2004; Chen et al. 2018; Patterson et al. 2019), and are parametrized by real numbers u and v . Here, we use the Dirac (bra-ket) notation to represent the quantum states as

$$|\psi_u\rangle = [\sqrt{1-u^2}, 0, u, 0]^\top, \quad (14)$$

$$|\psi_{v\pm}\rangle = [0, \pm\sqrt{1-v^2}, v, 0]^\top, \quad (15)$$

²<https://github.com/paddlepaddle/Quantum>

where $u, v \in [0, 1]$. Then, we can write this two sets of quantum states as a mixed quantum state ρ :

$$\rho(u, v) := q_1 \underbrace{|\psi_u\rangle\langle\psi_u|}_{\rho_1(u)} + \frac{q_2}{2} \underbrace{(|\psi_{v+}\rangle\langle\psi_{v+}| + |\psi_{v-}\rangle\langle\psi_{v-}|)}_{\rho_2(v)}, \quad (16)$$

with probability distribution $\{q_1 = \frac{1}{3}, q_2 = \frac{2}{3}\}$. These choices are consistent with the existing literature (Mohseni, Steinberg, and Bergou 2004; Chen et al. 2018).

Theoretical Distinguishability. We first analyze the ability of our method for classifying these two families of quantum states. The result is summarized in Theorem 6.

Theorem 6. *Given two families of non-orthogonal 2-qubit quantum states, shown in Eq. (16), and each has multiple copies. VSQL could exactly distinguish them by using only one shadow circuit, which consists of only one R_y rotation gate applied on a 1-local qubit.*

We defer the proof to Appendix. This Theorem shows that VSQL could theoretically distinguish these two different families of quantum states. We further evaluate the performance of VSQL via numerical experiments.

Experimental Setting and Results. 300 density matrices with 100 $\rho_1(u)$ (labeled as ‘0’) and 200 $\rho_2(v)$ (labeled as ‘1’) are sampled according to Eq. (16), where the parameters u and v are uniformly taken from $[0, 1]$. Then, we randomly select 80% of them as the training set and the rest 20% as the validation set. Consistent with the Theorem above, one 1-local shadow circuit, which consists of R_y gate only, is used to extract local features. The parameters of the shadow circuit θ and the FCNN $\{\mathbf{w}, b\}$ are initialized from the uniform distribution $\text{Uni}[0, 2\pi]$ and the Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbb{I})$, respectively. During the optimization loop, we choose the Adam (Kingma and Ba 2015) optimizer with a learning rate $\text{LR} = 0.03$. Learning curves for the training loss and the validation accuracy are illustrated in Fig. 3(a), where the distinguishability shown coincides with Theorem 6. We conclude VSQL could perfectly recognize the two families of quantum states defined in Eq. (16) after about 700 iterations. We find that the classification task becomes very difficult when u, v are both close to 1. This makes sense because $\rho_1(u = 1) \rightarrow \rho_2(v = 1)$ on the extreme case. This experimental result highlights the strength of our method. As a comparison, we adjust the sample range such that e.g., $u, v \in [0.1, 0.9]$, and the results are shown in Fig. 3(b). As expected, this modification leads to a faster convergence.

Classification of Multi-Class Quantum States. We also conduct the experiments for three-class case with 400 density matrices in total, where the data set is again taken from Eq. (16), but adding the third family $\rho_3(t) = |\psi_t\rangle\langle\psi_t|$ with a new probability distribution $\{q_i\} = \{\frac{1}{4}, \frac{1}{2}, \frac{1}{4}\}$, where

$$|\psi_t\rangle = [\sqrt{1-t^2}, t, 0, 0]^\top, \quad t \in [0, 1]. \quad (17)$$

The experimental settings are almost the same as the binary case and results are illustrated in Figs. 3(c) and 3(d). We conclude that VSQL also has a strong distinguishability for multi-class quantum states. A complete analysis about this block could be found in Appendix.

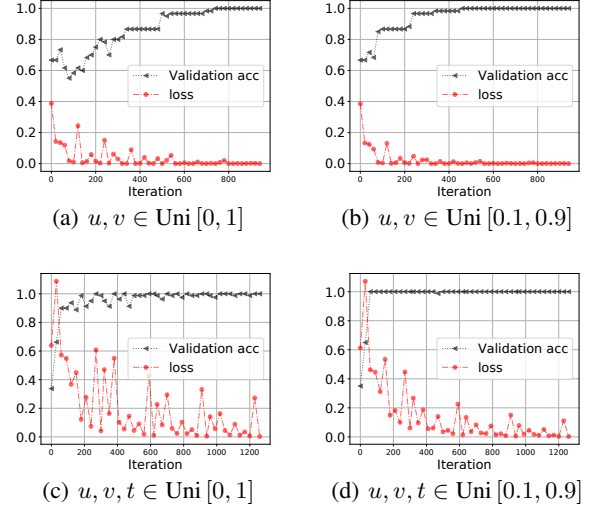


Figure 3: Learning curves that record the training loss and the validation accuracy of VSQL with different experimental settings. (a) and (b) are binary classification with different parameter range $u, v \in [0, 1]$ and $u, v \in [0.1, 0.9]$. By adjusting the sample range, the training loss and the validation accuracy reach the optimal values faster. (c) and (d) describe the similar experimental setting but for three-class classification of quantum states.

MNIST Classification

Next, we apply VSQL to classify handwritten digits taken from a public benchmark dataset MNIST (LeCun et al. 1998), which consists of 60,000 train examples and 10,000 test examples. The MNIST data set contains 10 different classes labeled from ‘0’ to ‘9’. Each image contains 28×28 gray scale pixels valued in $0 \sim 255$. In order to match the input of VSQL, these pictures are normalized and unfolded into 784-dimensional vectors. Then, we expand their dimension to 10-qubit pure quantum states $\{|\psi_i\rangle\}$ (1024-dimensional vectors) via zero-padding and represent them in the density matrix formulation $\{\rho_{in}^{(i)}\} = \{|\psi_i\rangle\langle\psi_i|\}$. By doing so, the pre-processing is complete and we obtain the training set $\mathcal{D}^{(train)} := \{(\rho_{in}^{(m)}, y^{(m)})\}_{m=1}^{N_{train}} \subset \mathbb{C}^{1024 \times 1024} \times \mathbb{R}^{10}$. We first select two classes (‘0’ and ‘1’) to verify the binary classification ability of VSQL, which contains 12,665 training samples (5923 0-label and 6742 1-label) and 2115 test samples (980 0-label and 1135 1-label). Then, we use the whole data set to evaluate the 10-class classification performance.

Experimental Setting. For the binary case, the 2-local shadow circuit (ansatz) used to extract local features is shown in Fig. 4. The number of repetitions of the dashed block structure is denoted as the circuit depth D and this ansatz has $2(D + 3)$ parameters in total. The parameters θ and $\{\mathbf{w}, b\}$ are initialized from a uniform distribution in $[0, 2\pi]$ and a Gaussian distribution $\mathcal{N}(\mathbf{0}, \mathbb{I})$, respectively. During the optimization, we choose the Adam optimizer




Figure 4: The 2-local shadow circuit design for MNIST classification (binary case). The first part uses $R_z - R_y - R_z$ combination to represent general rotations on each single-qubit subspace. The followed repeated block consists of CNOT gates and two single-qubit R_y rotations. The block circuit in the dashed box is repeated D times to extend the expressive power of quantum circuits.

Methods	n_s	D	# Ps	Test acc (%)
(Schuld et al. 2020)	/	/	124	96.70
(Farhi and Neven 2018)	/	/	96	98.00
VSQL (this paper)	1	1	18	99.43 ± 0.14
	2	1	35	99.52 ± 0.18

Table 1: Summary of the existing variational quantum classifiers on MNIST binary classification. VSQL outperforms other classifiers in terms of number of parameters and test accuracy by reaching 99.52% average test accuracy among 10 random experiments. # Ps denotes # Params.

with a batch size of 20 samples and a learning rate of $LR = 0.02$. Each experiment is repeated 10 times to collect the mean accuracy and the corresponding fluctuations. For the 10-class case, the classification task becomes much difficult and hence we choose 4-local shadow circuits to extract shadow features, which can be extended from the 2-local design in Fig. 4. There will be $4(D + 3)$ parameters in each shadow circuit. All the other settings are identical to the binary case, except for a new batch size of 200 samples.

Results. The results for the binary case are summarized in Table 1. Our method VSQL easily achieves an average test accuracy above 99% with only $n_s = 1$ shadow circuit and depth $D = 1$, which has 8 rotation angles in the shadow circuit and 9 weights and 1 bias in FCNN. This result demonstrates the powerful capacity of VSQL to classify handwritten digits. By adding another shadow circuit to $n_s = 2$ with 35 parameters, one could obtain an average test accuracy above 99.5%. As a comparison, we list the results of existing methods: Circuit-centric classifier (Schuld et al. 2020) and QNN classifier (Farhi and Neven 2018). Our method outperforms these variational quantum classifiers in terms of the number of parameters and test accuracy. Here, we should note that the details of their data preprocessing are slightly different from us, i.e., the Circuit-centric classifier uses the MNIST256 dataset with an 8-qubit quantum system, and the QNN classifier uses a 4×4 downsampled version of the MNIST dataset with a 17-qubit quantum system.

For multi-class classification, it is rarely discussed and tested in the literature of variational quantum classifiers. The one-vs.-all method is mentioned in Schuld et al. but troublesome to implement. Therefore, we only compare the perfor-

Methods	n_s	D	# Ps	Test acc (%)
NN (60k samples)	/	1	7850	91.63 ± 0.15
VSQL (60k samples)	5	5	520	87.69 ± 0.98
VSQL (60k samples)	9	5	928	91.13 ± 0.51
NN (1k samples)	/	1	7850	86.36 ± 0.23
VSQL (1k samples)	5	5	520	83.92 ± 1.20
VSQL (1k samples)	9	5	928	87.39 ± 0.40

Table 2: MNIST 10-class classification

mance of VSQL with a single-layered classical neural network (NN). The experimental settings of the classical neural network are similar to VSQL, and it contains 7840 weights and 10 biases to map the 784-dimension input vectors to 10-dimensional output vectors. The results are summarized in Table 2. When using 9 different shadow circuits with each circuit depth $D = 5$, VSQL could reach almost the same test accuracy with the single layer NN, but requiring much fewer parameters. Although this accuracy is not quite satisfied, it can still compete with the simplest classical NN. Notably, we find that if we select 1k samples (about 100 samples for each class) for training from 60k examples and choose the same size of test examples (10k), VSQL could achieve a higher test accuracy than NN (cf. the bottom half of Table 2). The above finding indicates that VSQL could extract high-level features from fewer training samples than NN, which may be a potential advantage of VSQL for future practical applications in the NISQ era.

Discussions

We propose the VSQL framework, which adopts the idea of classical shadows to distinguish quantum data. With theoretical justifications and numerical experiments, we have shown that VSQL for classification outperforms many other variational classifiers on the benchmark test of binary MNIST handwritten digit recognition with much fewer network parameters. In particular, in our framework, less noise will be introduced during the quantum-classical hybrid information processing as the number of quantum gates used is independent of the problem size. Simultaneously, by sampling a slice of the loss landscape, we briefly introduce the barren plateau problem and show the solution to escape from it. Notably, by adjusting the operating scope of shadow circuits, our approach can be easily implemented on the existing quantum devices with topological connectivity limits.

We believe that VSQL would open the possibility of many future directions. For example, it would be interesting to explore the applications on VSQL for generative models and unsupervised quantum machine learning tasks such as clustering. Furthermore, the online learning version of VSQL may also be a good future direction, see (Aaronson et al. 2018; Chen and Wang 2020; Yang et al. 2020). We also expect that VSQL may shed light on other quantum applications on near-term quantum devices.

Acknowledgements

We would like to thank Prof. Sanjiang Li, Prof. Yuan Feng and Youle Wang for helpful discussions. G. L. acknowledges the support from the Baidu-UTS AI Meets Quantum project and the China Scholarship Council (No. 201806070139). This work was partly supported by the Australian Research Council (Grant No: DP180100691). This work was done when Z. S. was a research intern at Baidu Research.

References

- Aaronson, S. 2018. Shadow tomography of quantum States. In *Proceedings of the Annual ACM Symposium on Theory of Computing*, 1088–1101. New York, New York, USA: ACM Press.
- Aaronson, S.; Chen, X.; Hazan, E.; and Kale, S. 2018. On-line learning of quantum states. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, 8976–8986.
- Bae, J.; and Kwek, L. C. 2015. Quantum state discrimination and its applications. *Journal of Physics A: Mathematical and Theoretical* 48(8): 083001.
- Benedetti, M.; Lloyd, E.; Sack, S.; and Fiorentini, M. 2019. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology* 4(4): 043001.
- Bhatia, A. S.; Saggi, M. K.; Kumar, A.; and Jain, S. 2019. Matrix Product State-Based Quantum Classifier. *Neural Computation* 31(7): 1499–1517.
- Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; and Lloyd, S. 2017. Quantum machine learning. *Nature* 549(7671): 195–202.
- Bottou, L. 2004. Stochastic Learning. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, volume 3176, 146–168.
- Cerezo, M.; Sharma, K.; Arrasmith, A.; and Coles, P. J. 2020a. Variational Quantum State Eigensolver. *arXiv:2004.01372* URL <http://arxiv.org/abs/2004.01372>.
- Cerezo, M.; Sone, A.; Volkoff, T.; Cincio, L.; and Coles, P. J. 2020b. Cost-Function-Dependent Barren Plateaus in Shallow Quantum Neural Networks. *arXiv:2001.00550* 1–35.
- Chen, H.; Wossnig, L.; Severini, S.; Neven, H.; and Mohseni, M. 2018. Universal discriminative quantum neural networks. *arXiv:1805.08654*.
- Chen, R.; Song, Z.; Zhao, X.; and Wang, X. 2020. Variational Quantum Algorithms for Trace Distance and Fidelity Estimation. *arXiv:2012.05768* 1–13. URL <http://arxiv.org/abs/2012.05768>.
- Chen, Y.; and Wang, X. 2020. More Practical and Adaptive Algorithms for Online Quantum State Learning. *arXiv preprint arXiv:2006.01013*.
- Ciliberto, C.; Herbster, M.; Ialongo, A. D.; Pontil, M.; Rocchetto, A.; Severini, S.; and Wossnig, L. 2018. Quantum machine learning: A classical perspective. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* 474(2209): 20170551.
- Dankert, C.; Cleve, R.; Emerson, J.; and Livine, E. 2009. Exact and approximate unitary 2-designs and their application to fidelity estimation. *Physical Review A* 80(1): 012304.
- Farhi, E.; Goldstone, J.; and Gutmann, S. 2014. A Quantum Approximate Optimization Algorithm. *arXiv:1411.4028* 1–16.
- Farhi, E.; and Harrow, A. W. 2016. Quantum Supremacy through the Quantum Approximate Optimization Algorithm. *arXiv:1602.07674* 1–22.
- Farhi, E.; and Neven, H. 2018. Classification with Quantum Neural Networks on Near Term Processors. *arXiv:1802.06002* 1–21.
- Gambs, S. 2008. Quantum classification. *arXiv:0809.0444* 119–123.
- Goodfellow, I.; Bengio, Y.; and Courville, A. 2016. *Deep Learning*. MIT Press. <http://www.deeplearningbook.org>.
- Grant, E.; Benedetti, M.; Cao, S.; Hallam, A.; Lockhart, J.; Stojevic, V.; Green, A. G.; and Severini, S. 2018. Hierarchical quantum classifiers. *npj Quantum Information* 4(1): 65.
- Harrow, A.; and Napp, J. 2019. Low-depth gradient measurements can improve convergence in variational hybrid quantum-classical algorithms. *arXiv:1901.05374* 1–45.
- Havlíček, V.; Córcoles, A. D.; Temme, K.; Harrow, A. W.; Kandala, A.; Chow, J. M.; and Gambetta, J. M. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature* 567(7747): 209–212.
- Huang, H.-Y.; Kueng, R.; and Preskill, J. 2020. Predicting many properties of a quantum system from very few measurements. *Nature Physics* 1–40.
- Kandala, A.; Mezzacapo, A.; Temme, K.; Takita, M.; Brink, M.; Chow, J. M.; and Gambetta, J. M. 2017. Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* 549(7671): 242–246.
- Kapoor, A.; Wiebe, N.; and Svore, K. 2016. Quantum perceptron models. In *Advances in Neural Information Processing Systems*, 3999–4007.
- Kingma, D. P.; and Ba, J. L. 2015. Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015 - Conference Track Proceedings*.
- Lecun, Y.; Bengio, Y.; and Hinton, G. 2015. Deep learning. *Nature* 521(7553): 436–444.
- LeCun, Y.; Bottou, L.; Bengio, Y.; and Haffner, P. 1998. Gradient-based learning applied to document recognition. *Proceedings of the IEEE* 86(11): 2278–2324.
- Li, T.; Chakrabarti, S.; and Wu, X. 2019. Sublinear quantum algorithms for training linear and kernel-based classifiers. In *International Conference on Machine Learning*, 3815–3824.

- Lloyd, S.; Mohseni, M.; and Rebentrost, P. 2013. Quantum algorithms for supervised and unsupervised machine learning. *arXiv:1307.0411* .
- Lloyd, S.; Schuld, M.; Ijaz, A.; Izaac, J.; and Killoran, N. 2020. Quantum embeddings for machine learning. *arXiv:2002.08953* 1–11.
- Ma, Y.; Yu, D.; Wu, T.; and Wang, H. 2019. PaddlePaddle: An Open-Source Deep Learning Platform from Industrial Practice. *Frontiers of Data and Computing* 1(1): 105–115.
- McClean, J. R.; Boixo, S.; Smelyanskiy, V. N.; Babbush, R.; and Neven, H. 2018. Barren plateaus in quantum neural network training landscapes. *Nature Communications* 9(1): 4812.
- Mitarai, K.; Negoro, M.; Kitagawa, M.; and Fujii, K. 2018. Quantum circuit learning. *Physical Review A* 98(3): 032309.
- Mohseni, M.; Steinberg, A. M.; and Bergou, J. A. 2004. Optical realization of optimal unambiguous discrimination for pure and mixed quantum states. *Physical Review Letters* 93(20): 200403.
- Nielsen, M. A.; and Chuang, I. 2002. Quantum Computation and Quantum Information. *American Journal of Physics* 70(5): 558–559.
- Patterson, A.; Chen, H.; Wossnig, L.; Severini, S.; Browne, D.; and Rungger, I. 2019. Quantum State Discrimination Using Noisy Quantum Neural Networks. *arXiv:1911.00352* .
- Peruzzo, A.; McClean, J.; Shadbolt, P.; Yung, M.-H.; Zhou, X.-Q.; Love, P. J.; Aspuru-Guzik, A.; and O’Brien, J. L. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* 5(1): 4213. ISSN 2041-1723. doi:10.1038/ncomms5213. URL <http://www.nature.com/articles/ncomms5213>.
- Preskill, J. 2018. Quantum Computing in the NISQ era and beyond. *Quantum* 2: 79.
- Roa, L.; Retamal, J. C.; and Saavedra, C. 2002. Quantum-state discrimination. *Physical Review A - Atomic, Molecular, and Optical Physics* 66(1): 121031–121034.
- Schuld, M.; Bocharov, A.; Svore, K. M.; and Wiebe, N. 2020. Circuit-centric quantum classifiers. *Physical Review A* 101(3): 032308.
- Schuld, M.; and Killoran, N. 2019. Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters* 122(4): 040504.
- Schuld, M.; and Petruccione, F. 2018. Quantum ensembles of quantum classifiers. *Scientific Reports* 8(1): 2772.
- Schuld, M.; Sinayskiy, I.; and Petruccione, F. 2015. An introduction to quantum machine learning. *Contemporary Physics* 56(2): 172–185.
- Wang, S.; Fontana, E.; Cerezo, M.; Sharma, K.; Sone, A.; Cincio, L.; and Coles, P. J. 2020. Noise-Induced Barren Plateaus in Variational Quantum Algorithms. *arXiv:2007.14384* 2(Theorem 1): 1–20.
- Wang, X.; Song, Z.; and Wang, Y. 2020. Variational Quantum Singular Value Decomposition. *arXiv:2006.02336* URL <http://arxiv.org/abs/2006.02336>.
- Yang, F.; Jiang, J.; Zhang, J.; and Sun, X. 2020. Revisiting Online Quantum State Learning. In *AAAI*, 6607–6614.
- Yuan, X.; Endo, S.; Zhao, Q.; Li, Y.; and Benjamin, S. C. 2019. Theory of variational quantum simulation. *Quantum* 3: 191. ISSN 2521-327X. doi:10.22331/q-2019-10-07-191. URL <https://quantum-journal.org/papers/q-2019-10-07-191/>.
- Ziegel, E. R.; Lehmann, E. L.; and Casella, G. 1999. Theory of Point Estimation. *Technometrics* 41(3): 274.