



OPEN Variational quantum approximate support vector machine with inference transfer

Siheon Park¹, Daniel K. Park^{2,3} & June-Koo Kevin Rhee^{1,4}✉

A kernel-based quantum classifier is the most practical and influential quantum machine learning technique for the hyper-linear classification of complex data. We propose a Variational Quantum Approximate Support Vector Machine (VQASVM) algorithm that demonstrates empirical sub-quadratic run-time complexity with quantum operations feasible even in NISQ computers. We experimented our algorithm with toy example dataset on cloud-based NISQ machines as a proof of concept. We also numerically investigated its performance on the standard Iris flower and MNIST datasets to confirm the practicality and scalability.

Quantum computing opens up new exciting prospects of quantum advantages in machine learning in terms of sample and computation complexity^{1–5}. One of the foundations of these quantum advantages is the ability to form and manipulate data efficiently in a large quantum feature space, especially with kernel functions used in classification and other classes of machine learning^{6–14}.

The support vector machine (henceforth SVM)¹⁵ is one of the most comprehensive models that help conceptualize the basis of supervised machine learning. SVM classifies data by finding the optimal hyperplane associated with the widest margin between the two classes in a feature space. SVM can also perform highly nonlinear classifications using what is known as the kernel trick^{16–18}. The convexity of SVM guarantees global optimization.

One of the first quantum algorithms exhibiting an exponential speed-up capability is the least-square quantum support vector machine (LS-QSVM)⁵. However, the quantum advantage of LS-QSVM strongly depends on costly quantum subroutines such as density matrix exponentiation¹⁹ and quantum matrix inversion^{20,21} as well as components such as quantum random access memory (QRAM)^{2,22}. Because the corresponding procedures require quantum computers to be fault-tolerant, LS-QSVM is unlikely to be realized in noisy intermediate-scale quantum (NISQ) devices²³. On the other hand, there are a few quantum kernel-based machine-learning algorithms for near-term quantum applications. Well-known examples are quantum kernel estimators (QKE)⁸, variational quantum classifiers (VQC)⁸, and Hadamard or SWAP test classifiers (HTC, STC)^{10,11}. These algorithms are applicable to NISQ, as there are no costly operations needed. However, the training time complexity is even worse than in the classical SVM case. For example, the number of measurements required to generate only the kernel matrix evaluation of QKE scales with the number of training samples to the power of four⁸.

Here, we propose a novel quantum kernel-based classifier that is feasible with NISQ devices and that can exhibit a quantum advantage in terms of accuracy and training time complexity as exerted in Ref.⁸. Specifically, we have discovered distinctive designs of quantum circuits that can evaluate the objective and decision functions of SVM. The number of measurements for these circuits with a bounded error is independent from the number of training samples. The depth of these circuits scales also linearly with the size of the training dataset. Meanwhile, the exponentially fewer parameters of parameterized quantum circuits (PQCs)²⁴ encodes the Lagrange multipliers of SVM. Therefore, the training time of our model with a variational quantum algorithm (VQA)²⁵ scales as sub-quadratic, which is asymptotically lower than that of the classical SVM case^{5,26,27}. Our model also shows an advantage in classification due to its compatibility with any typical quantum feature map.

Results

Support vector machine (SVM). Data classification infers the most likely class of an unseen data point $\hat{x} \in \mathbb{C}^N$ given a training dataset $\mathcal{S} = \{(\mathbf{x}_i, y_i)\}_{i=0}^{M-1} \subset \mathcal{X} \times \mathcal{Y}$. Here, $\mathcal{X} \subset \mathbb{C}^N$ and $\mathcal{Y} = \{0, 1, \dots, L-1\}$. Although the data is real-valued in practical machine learning tasks, we allow complex-valued data without a loss of generality. We focus on binary classification. (i.e., $\mathcal{Y} = \{-1, 1\}$), because multi-class classification can be

¹KAIST, School of Electrical Engineering, Daejeon 34141, South Korea. ²Department of Applied Statistics, Yonsei University, Seoul 03722, South Korea. ³Department of Statistics and Data Science, Yonsei University, Seoul 03722, South Korea. ⁴Qunova Computing, Inc., Daejeon 34051, South Korea. ✉email: rhee.jk@kaist.edu

conducted with a multiple binary SVM via a one-versus-all or a one-versus-on scheme²⁸. We assume that \mathcal{S} is linearly separable in the higher dimensional Hilbert space \mathcal{H} given some feature map $\phi : \mathcal{X} \mapsto \mathcal{H}$. Then, there should exist two parallel supporting hyperplanes $(\mathbf{w}, \phi(\cdot)) + b = y \in \mathcal{Y}$ that divide training data. The goal is to find hyperplanes for which the margin between them is maximized. To maximize the margin even further, the linearly separable condition can be relaxed so that some of the training data can penetrate into the “soft” margin. Because the margin is given as $2/\|\mathbf{w}\|$ by simple geometry, the mathematical formulation of SVM¹³ is given as

$$p^* = \min_{\mathbf{w}, b, \xi} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{C}{2} \sum_{i=0}^{M-1} \xi_i^2 : y_i((\mathbf{w}, \phi(\mathbf{x}_i)) + b) \geq 1 - \xi_i, \quad (1)$$

where the slack variable ξ is introduced to represent a violation of the data in the linearly separable condition. The dual formulation of SVM is expressed as²⁹

$$d^* = \max_{\beta \geq 0} \sum_{i=0}^{M-1} \beta_i - \frac{1}{2} \sum_{i,j=0}^{M-1} \beta_i \beta_j y_i y_j k(\mathbf{x}_i, \mathbf{x}_j) - \frac{1}{2C} \sum_{i=0}^{M-1} \beta_i^2 : \sum_{i=0}^{M-1} \beta_i y_i = 0, \quad (2)$$

where the positive semi-definite (PSD) kernel is $k(\mathbf{x}_1, \mathbf{x}_2) = \langle \phi(\mathbf{x}_1), \phi(\mathbf{x}_2) \rangle$ for $\mathbf{x}_{1,2} \in \mathcal{X}$. The β_i values are non-negative Karush-Kuhn-Tucker multipliers. This formulation employs an implicit feature map uniquely determined by the kernel. The global solution β^* is obtained in polynomial time due to convexity²⁹. After optimization, the optimum bias is recovered as $b^* = y_q(1 - C^{-1} \beta_q^*) - \sum_{i=0}^{M-1} \beta_i^* y_i k(\mathbf{x}_q, \mathbf{x}_i)$ for any $\beta_q^* > 0$. Such training data \mathbf{x}_q with non-zero weight β_q are known as the support vectors. We estimate the labels of unseen data with a binary classifier:

$$\hat{y} = \text{sgn} \left\{ \sum_{i=0}^{M-1} \beta_i^* y_i k(\mathbf{x}_i, \hat{\mathbf{x}}) + b^* \right\}. \quad (3)$$

In a first-hand principle analysis, the complexity of solving Eq. (2) is $\mathcal{O}(M^2(N + M) \log(1/\delta))$ with accuracy of δ . A kernel function with complexity of $\mathcal{O}(N)$ is queried $M(M - 1)/2$ times to construct the kernel matrix, and quadratic programming takes $\mathcal{O}(M^3 \log(1/\delta))$ to find β^* for a non-sparse kernel matrix⁵. Although the complexity of SVM decreases when employing modern programming methods^{26,27}, it is still higher than or equal to $\mathcal{O}(M^2N)$ due to kernel matrix generation and quadratic programming. Thus, a quantum algorithm that evaluates all terms in Eq. (2) for $\mathcal{O}(MN)$ time and achieves a minimum of fewer than $\mathcal{O}(M)$ evaluations would have lower complexity than classical algorithms. We apply two forms of transformations to Eq. (2) to realize an efficient quantum algorithm.

Change of variable and bias regularization. Constrained programming, such as that in the SVM case, is often transformed into unconstrained programming by adding penalty terms of constraints to the objective function. Although there are well-known methods such as an interior point method²⁹, we prefer the strategies of a ‘change of variables’ and ‘bias regularization’ to maintain the quadratic form of SVM. Although motivated to eliminate constraints, the results appear likely to lead to an efficient quantum SVM algorithm.

First, we change optimization variable β to (α, B) , where $B := \sum_{i=0}^{M-1} \beta_i$ and $\alpha := \beta/B$ to eliminate inequality constraints. The l_1 -normalized variable α is an M -dimensional probability vector given that $0 \leq \alpha_i \leq 1, \forall i \in \{0, \dots, M-1\}$ and $\sum_{i=0}^{M-1} \alpha_i = 1$. Let us define $W_k(\alpha; \mathcal{S}) := \sum_{i,j=0}^{M-1} \alpha_i \alpha_j y_i y_j k(\mathbf{x}_i, \mathbf{x}_j) + C^{-1} \sum_{i=0}^{M-1} \alpha_i^2$. We substitute the variables into Eq. (2):

$$\max_{\alpha \in PV_M} \max_{B \geq 0} \left\{ B - \frac{1}{2} B^2 W_k(\alpha; \mathcal{S}) \right\} : \sum_{i=0}^{M-1} \alpha_i y_i = 0, \quad (4)$$

where PV_M is a set of M -dimensional probability vectors. Because $W_k(\alpha; \mathcal{S}) \geq 0$ for an arbitrary α due to the property of the positive semi-definite kernel, $B^* = 1/W_k(\alpha; \mathcal{S})$ is a partial solution that maximizes Eq. (4) on B . Substituting B^* with Eq. (4), we have

$$\max_{\alpha \in PV_M} \frac{1}{2W_k(\alpha; \mathcal{S})} : \sum_{i=0}^{M-1} \alpha_i y_i = 0. \quad (5)$$

Finally, because maximizing $1/2W_k(\alpha; \mathcal{S})$ is identical to minimizing $W_k(\alpha; \mathcal{S})$, we have a simpler formula that is equivalent to Eq. (2):

$$\tilde{d}^* = \min_{\alpha \in PV_M} \sum_{i,j=0}^{M-1} \alpha_i \alpha_j y_i y_j k(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{C} \sum_{i=0}^{M-1} \alpha_i^2 : \sum_{i=0}^{M-1} \alpha_i y_i = 0, \quad (6)$$

The above Eq. (6) implies that instead of optimizing M numbers of bounded free parameters β or α , we can optimize the $\log(M)$ -qubit quantum state $|\psi_\alpha\rangle$ and define $\alpha_i := |\langle i | \psi_\alpha \rangle|^2$. Therefore, if there exists an efficient quantum algorithm that evaluates the objective function of Eq. (6) given $|\psi_\alpha\rangle$, the complexity of SVM would be improved. In fact, in the later section, we propose quantum circuits with linearly scaling complexity for that purpose.

The equality constraint is relaxed after adding the l_2 -regularization term of the bias to Eq. (1). Motivated by the loss function and regularization perspectives of SVM³⁰, this technique was introduced^{31,32} and developed^{33,34} previously. The primal and dual forms of SVM become

$$p^* = \min_{\mathbf{w}, b, \xi} \frac{1}{2} \|\mathbf{w}\|^2 + \frac{\lambda}{2} b^2 + \frac{C}{2} \sum_{i=0}^{M-1} \xi_i^2 : y_i(\langle \mathbf{w}, \phi(\mathbf{x}_i) \rangle + b) \geq 1 - \xi_i, \tag{7}$$

$$d^* = \max_{\beta \geq 0} \sum_{i=0}^{M-1} \beta_i - \frac{1}{2} \sum_{i,j=0}^{M-1} \beta_i \beta_j y_i y_j \left[k(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{\lambda} \right] - \frac{1}{2C} \sum_{i=0}^{M-1} \beta_i^2. \tag{8}$$

Note that $k(\cdot, \cdot) + \lambda^{-1}$ is a positive definite. As shown earlier, changing the variables causes Eq. (8) to become another equivalent optimization problem:

$$\tilde{d}^* = \min_{\alpha \in PV_M} \sum_{i,j=0}^{M-1} \alpha_i \alpha_j y_i y_j \left[k(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{\lambda} \right] + \frac{1}{C} \sum_{i=0}^{M-1} \alpha_i^2. \tag{9}$$

As the optimal bias is given as $b^* = \lambda^{-1} \sum_{i=0}^{M-1} \alpha_i^* y_i$ according to the Karush-Kuhn-Tucker condition, the classification formula inherited from Eq. (3) is expressed as

$$\hat{y} = \text{sgn} \left\{ \sum_{i=0}^{M-1} \alpha_i^* y_i \left[k(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{\lambda} \right] \right\}. \tag{10}$$

Equations (8) and (9) can be viewed as Eqs. (2) and (6) with a quadratic penalizing term on the equality constraint such that they become equivalent in terms of the limit of $\lambda \rightarrow 0$. Thus, Eqs. (7), (8), and (9) are more relaxed SVM optimization problems with an additional hyperparameter λ .

Variational quantum approximate support vector machine. One way to generate the aforementioned quantum state $|\psi_\alpha\rangle$ is to use amplitude encoding: $|\psi_\alpha\rangle = \sum_{i=0}^{M-1} \sqrt{\alpha_i} |i\rangle$. However, doing so would be inefficient because the unitary gate of amplitude encoding has a complex structure that scales as $\mathcal{O}(\text{poly}(M))$ ³⁵. Another way to generate $|\psi_\alpha\rangle$ is to use a parameterized quantum circuit (PQC), known as an ansatz in this case. Because there is no prior knowledge in the distribution of α_i^* s, the initial states should be $|+\dots+\rangle = \frac{1}{\sqrt{M}} \sum_{i=0}^{M-1} |i\rangle$. The ansatz $V(\theta)$ can transform the initial state into other states depending on gate parameter vector θ : $|\psi_\alpha\rangle = V(\theta)|+\dots+\rangle$. In other words, optimization parameters encoded by θ with the ansatz are represented as $\alpha_i(\theta) = |\langle i|V(\theta)|+\dots+\rangle|^2$. Given the lack of prior information, the most efficient ansatz design can be a hardware-efficient ansatz (HEA), which consists of alternating local rotation layers and entanglement layers^{36,37}. The number of qubits and the depth of this ansatz are $\mathcal{O}(\text{polylog}(M))$.

We discovered the quantum circuit designs that compute Eqs. (9) and (10) within $\mathcal{O}(M)$ time. Conventionally, the quantum kernel function is defined as the Hilbert Schmidt inner product: $k(\cdot, \cdot) = |\langle \phi(\cdot) | \phi(\cdot) \rangle|$ ^{24,7,8,10-12}. First, we divide the objective function in Eq. (9) into the loss and regularizing functions of θ using the above ansatz encoding:

$$\mathcal{L}_{\phi,\lambda}(\theta; \mathcal{S}) = \sum_{i,j=0}^{M-1} \alpha_i(\theta) \alpha_j(\theta) y_i y_j \left[|\langle \phi(\mathbf{x}_i) | \phi(\mathbf{x}_j) \rangle|^2 + \frac{1}{\lambda} \right], \mathcal{R}(\theta) = \sum_{i=0}^{M-1} \alpha_i(\theta)^2. \tag{11}$$

Specifically, the objective function is equal to $\mathcal{L}_{\phi,\lambda} + C^{-1}\mathcal{R}$. Similarly, the decision function in Eq. (10) becomes

$$f_{\phi,\lambda}(\mathbf{x}; \theta, \mathcal{S}) = \sum_{i=0}^{M-1} \alpha_i(\theta) y_i \left[|\langle \phi(\mathbf{x}_i) | \phi(\mathbf{x}) \rangle|^2 + \frac{1}{\lambda} \right]. \tag{12}$$

Inspired by STC^{10,11}, the quantum circuits in Fig. 1 efficiently evaluate $\mathcal{L}_{\phi,\lambda}, \mathcal{R}$ and $f_{\phi,\lambda}$. The quantum gate $\mathcal{U}_{\phi,\mathcal{S}}$ embeds the entire training dataset with the corresponding quantum feature map $U_{\phi(\mathbf{x})}|00\dots0\rangle = |\phi(\mathbf{x})\rangle$, so that $\mathcal{U}_{\phi,\mathcal{S}}|i\rangle \otimes |00\dots0\rangle \otimes |0\rangle = |i\rangle \otimes |\phi(\mathbf{x}_i)\rangle \otimes |y_i\rangle$. Therefore, the quantum state after state preparation is $|\Psi\rangle = \sum_{i=0}^{M-1} \sqrt{\alpha_i} |i\rangle \otimes |\phi(\mathbf{x}_i)\rangle \otimes |y_i\rangle$. We apply a SWAP test and a joint σ_z measurement in the loss and decision circuits to evaluate $\mathcal{L}_{\phi,\lambda}$ and $f_{\phi,\lambda}$:

$$\mathcal{L}_{\phi,\lambda}(\theta; \mathcal{S}) = \langle \sigma_z^a \sigma_z^{y_1} \sigma_z^{y_2} \rangle_\theta + \frac{1}{\lambda} \langle \sigma_z^{y_1} \sigma_z^{y_2} \rangle_\theta, f_{\phi,\lambda}(\mathbf{x}; \theta, \mathcal{S}) = \langle \sigma_z^a \sigma_z^y \rangle_{\mathbf{x};\theta} + \frac{1}{\lambda} \langle \sigma_z^y \rangle_{\mathbf{x};\theta}, \mathcal{R}(\theta) = \langle M_{00\dots0} \rangle_\theta. \tag{13}$$

Here, σ_z is a Pauli Z operator and $M_{00\dots0}$ is a projection measurement operator of state $|0\rangle^{\otimes \log(M)}$. See “Methods” for specified derivations. The asymptotic complexities of the loss and decision circuits are linear with regard to the amount of training data. $\mathcal{U}_{\phi,\mathcal{S}}$ can be prepared with $\mathcal{O}(MN)$ operations⁴. See “Methods” for the specific realization used in this article. Because $V(\theta)$ has $\mathcal{O}(\text{polylog}(M))$ depth and the SWAP test requires only $\mathcal{O}(N)$ operations, the overall run-time complexity of evaluating $\mathcal{L}_{\phi,\lambda}(\theta; \mathcal{S})$ and $f_{\phi,\lambda}(\mathbf{x}; \theta, \mathcal{S})$ with bounded error ϵ

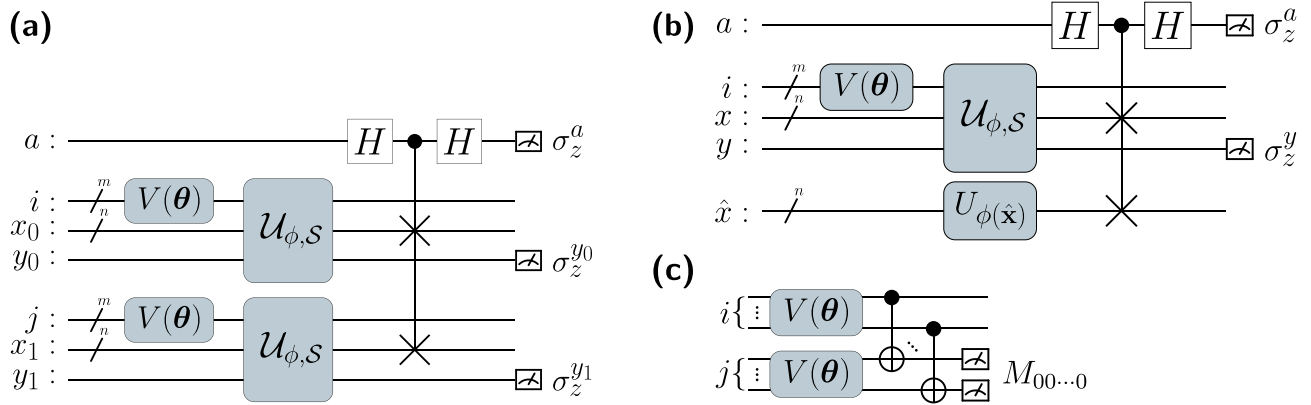


Figure 1. Circuit architecture of VQASVM. Loss, decision, and regularization circuits are shown in the order of panel (a,b), and (c) all qubits of index registers i and j are initialized to $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$, and the rest to $|0\rangle$. Ansatz $V(\theta)$ is a PQC of $m = \log(M)$ qubits that encodes probability vector α . $U_{\phi,S}$ embeds a training data set S with a quantum feature map $U_{\phi(\hat{x})}$, which embeds classical data \hat{x} to a quantum state $|\phi(\hat{x})\rangle$. n denotes the number of qubits for the quantum feature map, which is usually N , but can be reduced to $\log(N)$ if an amplitude encoding feature map is used.

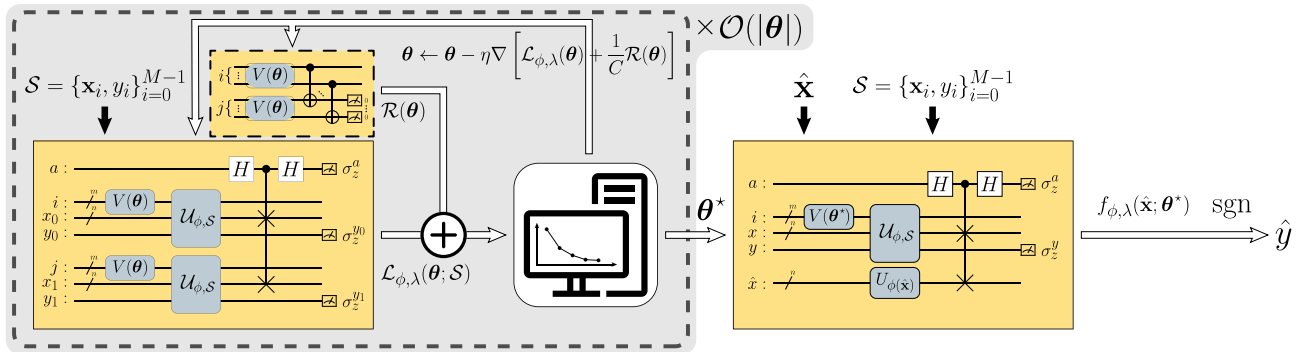


Figure 2. Variational quantum approximated support vector machine. The white round boxes represent classical calculations whereas the yellow round boxes represent quantum operations. The white arrows represent the flow of classical data whereas the black arrows represent the embedding of classical data. The grey areas indicate the corresponding training phase of each iteration. The regularization circuit in the black dashed box can be omitted for a hard-margin case where $C \rightarrow \infty$.

is $\mathcal{O}(\varepsilon^{-2}MN)$ (see the Supplementary Fig. S5 online for numerical verification of the scaling). Similarly, the complexity of estimating $\mathcal{R}(\theta)$ with accuracy ε is $\mathcal{O}(\varepsilon^{-2}\text{polylog}(M))$ due to two parallel $V(\theta)$ s and $\mathcal{O}(\log(M))$ CNOT gates.

We propose a variational quantum approximate support vector machine (VQASVM) algorithm that solves the SVM optimization problem with VQA²⁵ and transfers the optimized parameters to classify new data effectively. Figure 2 summarizes the process of VQASVM. We estimate θ^* , which minimizes the objective function; this is then used for classifying unseen data:

$$\theta^* = \arg \min_{\theta} \mathcal{L}_{\phi,\lambda}(\theta; S) + \frac{1}{C} \mathcal{R}(\theta), \hat{y} = \text{sgn}\{f_{\phi,\lambda}(\hat{x}; \theta^*, S)\}. \tag{14}$$

Following the general scheme of VQA, the gradient descent (GD) algorithm can be applied; classical processors update the parameters of θ_i , whereas the quantum processors evaluate the functions for computing gradients. Because the objective function of VQASVM can be expressed as the expectation value of a Hamiltonian, i.e.,

$$\mathcal{L}(\theta; S) + \frac{1}{C} \mathcal{R}(\theta) = \langle +|V(\theta)^\dagger \otimes V(\theta)^\dagger H V(\theta) \otimes V(\theta)|+\rangle, \tag{15}$$

where $H = \sum_{ij=0}^{M-1} [y_i y_j k(\mathbf{x}_i, \mathbf{x}_j) + \frac{1}{\lambda} y_i y_j + \frac{1}{C} \delta_{ij}] |i\rangle\langle i| \otimes |j\rangle\langle j|$, the exact gradient can be obtained by the modified parameter-shift rule^{38,39}. GD converges to a local minimum after $\mathcal{O}(\log(1/\delta))$ iterations with the difference δ ²⁹ given that estimation error of the objective function is smaller than δ . Therefore, the total run-time complexity of VQASVM is $\mathcal{O}(\varepsilon^{-2} \log(1/\varepsilon) MN \text{polylog}(M))$ with error of ε as the number of parameters is $\mathcal{O}(\text{polylog}(M))$.

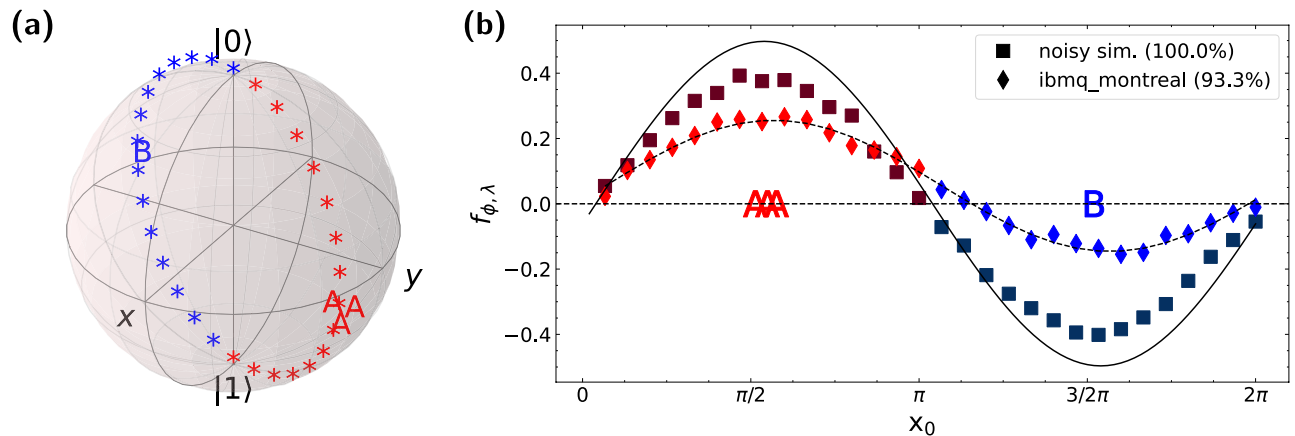


Figure 3. Experiments on a *ibmq_montreal* cloud NISQ processor. **(a)** The toy training (letters) and test (asterisk) data are shown here in a Bloch sphere. The color indicates the true class label of the data; i.e., red = class A, and blue = class B. The letters A represent the training data of class A, and letter B represents the training datum of class B. **(b)** classification results performed on *ibmq_montreal* QPU (diamonds) and a simulation with noise (squares) compared to theoretical values (solid line). $f_{\phi, \lambda}$ is the decision function value of each test datum. The letters A and B represent the training data, located at their longitudinal coordinates on the Bloch sphere (x_0). Curved dashed lines are the sine-fitting of the *ibmq_montreal* results. Values inside the round brackets in the legend are the classification accuracy rates.

Experiments on IBM quantum processors. We demonstrate the classification of a toy dataset using the VQASVM algorithm on NISQ computers as a proof-of-concept. Our example dataset is mapped to a Bloch sphere, as shown in Fig. 3a. Due to decoherence, we set the data dimension to $N = 2$ and number of training data instances to $M = 4$. First, we randomly choose the greatest circle on the Bloch sphere that passes $|0\rangle$ and $|1\rangle$. Then, we randomly choose two opposite points on the circle to be the center of two classes, A and B. Subsequently, four training data instances are generated close to each class center in order to avoid overlaps between the test data and each other. This results in a good training dataset with the maximum margin such that soft-margin consideration is not needed. In addition, thirty test data instances are generated evenly along the great circle and are labelled as 1 or -1 according to the inner products with the class centers. In this case, we can set hyperparameter $C \rightarrow \infty$ and the process hence requires no regularization circuit evaluation. The test dataset is non-trivial to classify given that the test data are located mostly in the margin area; convex hulls of both training datasets do not include most of the test data.

We choose a quantum feature map that embeds data (x_0, x_1) into a Bloch sphere instead of $N = 2$ qubits: $U_{\phi(x_0, x_1)} = R_z(x_1)R_y(x_0)$. Features x_0 and x_1 are the latitude and the longitude of the Bloch sphere. We use two qubits (q_0 and q_1) *RealAmplitude*⁴⁰ PQC as the ansatz: $V(\theta) = R_y^{q_0}(\theta_2) \otimes R_y^{q_1}(\theta_3)CNOT_{q_0 \rightarrow q_1}R_y^{q_0}(\theta_0) \otimes R_y^{q_1}(\theta_1)$. In this experiment, we use *ibmq_montreal*, which is one of the IBM Quantum Falcon processors. “Methods” section presents the specific techniques for optimizing quantum circuits against decoherence. The simultaneous perturbation stochastic approximation (SPSA) algorithm is selected to optimize $V(\theta)$ due to its rapid convergence and good robustness to noise^{41,42}. The measurements of each circuits are repeated $R = 8192$ times to estimate expectation values, which was the maximum possible option for *ibmq_montreal*. Due to long queue time of a cloud-based QPU, we reduce the QPU usage by applying warm-start and early-stopping techniques explained in “Methods” section.

Figure 3 shows the classification result. Theoretical decision function values were calculated by solving Eq. (9) with convex optimization. A noisy simulation is a classical simulation that emulates an actual QPU based on a noise parameter set estimated from a noise measurement. Although the scale of the decision values is reduced, the impact on the classification accuracy is negligible given that only the signs of decision values matter. This logic has been applied to most NISQ-applicable quantum binary classifiers. The accuracy would improve with additional error mitigation and offset calibration processes on the quantum device. Other VQASVM demonstrations with different datasets can be found as Supplementary Fig. S2 online.

Numerical simulation. In a practical situation, the measurement on quantum circuits is repeated R times to estimate the expectation value within $\varepsilon = \mathcal{O}(1/\sqrt{R})$ error, which could interfere with the convergence of VQA. However, the numerical analysis with the Iris dataset⁴³ confirmed that VQASVM converges even with the noise in objective function estimation exists. The following paragraphs describe the details of the numerical simulation, such as data preprocessing and the choice of the quantum kernel and the ansatz.

We assigned the labels $+1$ to Iris setosa and -1 to Iris versicolour and Iris virginica for binary classification. The features of the data were scaled so that the range became $[-\pi, \pi]$. We sampled $M = 64$ training data instances from the total dataset and treated the rest as the test data. The training kernel matrix constructed with our custom quantum feature map in Fig. 4a is learnable; i.e., the singular values of the kernel matrix decay exponentially. After testing the PQC designs introduced in Ref.³⁶, we chose the PQC exhibited in Fig. 4b as the ansatz for this

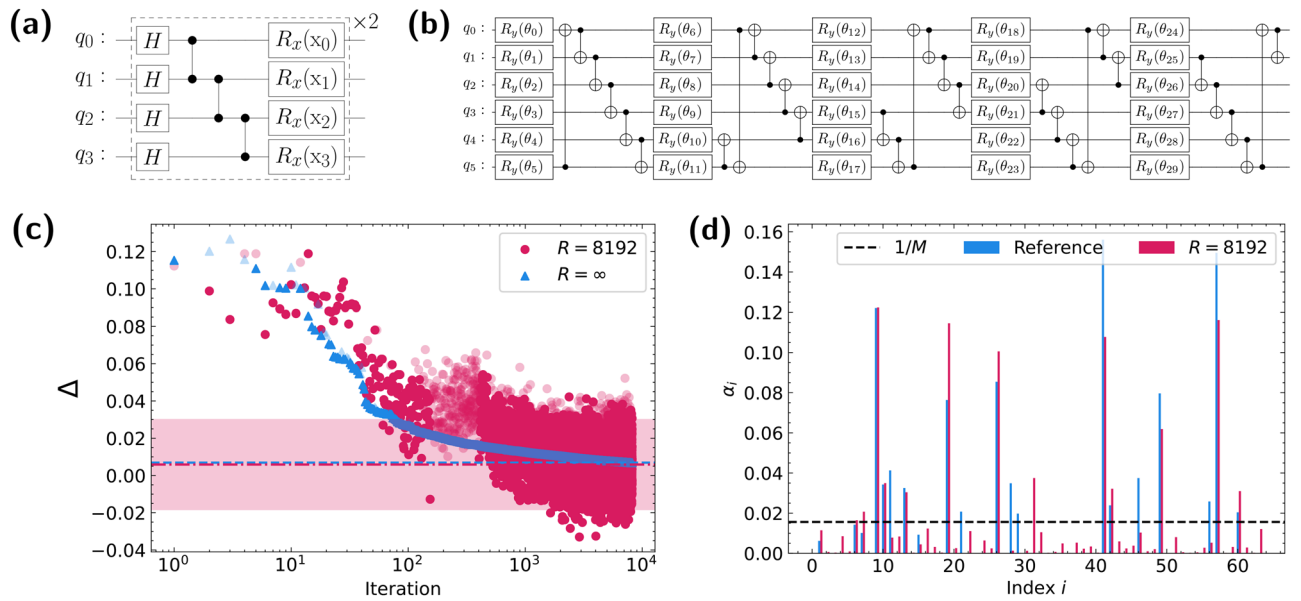


Figure 4. Numerical analysis on the iris dataset ($\lambda = C = 10^4$). **(a)** Custom quantum feature map for the iris dataset. **(b)** PQC design for $V(\theta)$ with $5 \times \log(M) = 30$ gate parameters. **(c)** The shaded circles and triangles depict the training convergence outcomes of the residual losses. At the final iteration, ($t = 2^{13}$) the residual loss Δ for $R = 8192$ repeated measurements (red dot-dashed line) is almost equal to the case of $R = \infty$ (blue dashed line), where the error when estimating the expectation value is 0. The red shaded area represents the 95% credible intervals of the last 16 residual losses for the $R = 8192$ case. **(d)** The spectrum of optimized weights, α_i s, of \mathbf{x}_i for the theoretical and $R = 8192$ cases are compared. The dashed black line indicates the level of the uniform weight $\alpha_i = 1/M = 1/64$. ‘Reference’ in the legend refers to the theoretical values of the α_i s obtained by convex optimization.

simulation (see Supplementary Figs. S6–S9 online). The number of PQC parameters is 30, which is less than M . In this simulation, the SPSA optimizer was used for training due to its fast convergence and robustness to noise^{41,42}.

The objective and decision functions were evaluated in two scenarios. The first case samples a finite number of measurement results $R = 8192$ to estimate the expectation values such that the error of the estimation is non-zero. The second case directly calculates the expectation values with zero estimation error, which corresponds to sampling infinitely many measurement results; i.e., $R = \infty$. We defined the residual loss of training as $\Delta = \mathcal{L}_{\phi, \lambda}(\theta^t; \mathcal{S}) + C^{-1} \mathcal{R}(\theta^t) - \tilde{d}^*$ at iteration t to compare the convergence. Here, \tilde{d}^* is the theoretical minimum of Eq. (9) as obtained by convex optimization.

Although containing some uncertainty, Fig. 4c shows that SPSA converges to a local minimum despite the estimation noise. Both the $R = 8192$ and $R = \infty$ cases show the critical convergence rule of SPSA; $\Delta \sim \mathcal{O}(|\theta|/t)$ for a sufficiently large number of iterations t . More vivid visualization can be found as Supplementary Fig. S10 online. In addition, the spectrum of the optimized Lagrange multipliers α_i s mostly coincides with the theory, especially for the significant support vectors. (Fig. 4d) Therefore, we concluded that training VQASVM within a finite number of measurements is achievable. The classification accuracy was 95.34% for $R = 8192$ and 94.19% for $R = \infty$.

The empirical speed-up of VQASVM is possible because the number of optimization parameters is exponentially reduced by the PQC encoding. However, in the limit of large number of qubits (i.e., $\log(M)$), it is unclear that such PQC can be trained with VQA to well-approximate the optimal solution. Therefore, we performed numerical analysis to empirically verify that VQASVM with $\mathcal{O}(\text{polylog}(M))$ parameters achieves bounded classification error even for large M . For this simulation, the MNIST dataset⁴⁴ was used instead of the Iris dataset because there are not enough Iris data points to clarify the asymptotic behavior of VQASVM. In this setting, the number of maximum possible MNIST training data instances for VQASVM is $M = 2^{13}$.

A binary image data of ‘0’s and ‘1’s with a 28×28 image size were selected for binary classification, the features of which were then reduced to 10 by means of a principle component analysis (PCA). The well-known quantum feature map introduced in Ref.⁸ was chosen for the simulation: $U_{\phi(\mathbf{x})} = W_{\phi(\mathbf{x})} H^{\otimes n} W_{\phi(\mathbf{x})} H^{\otimes n}$, where $W_{\phi(\mathbf{x})} = \exp i \sum_{G \subset [n], |G|=2} g_G(\mathbf{x}) \prod_{i \in G} \sigma_z^i$ and $g_{\{i\}}(\mathbf{x}) = x_i, g_{\{i,j\}}(\mathbf{x}) = (\pi - x_i)(\pi - x_j) \delta_{i+1,j}$. The visualization of the feature map can be found as Supplementary Fig. S4 online. The ansatz architecture used for this simulation was the PQC template shown in Fig. 4b with 19 layers; i.e., the first part of the PQC in Fig. 4b is repeated 19 times. Thus, the number of optimization parameters is $19 \times \log(M)$.

The numerical simulation shows that although residual training loss Δ linearly increases with the number of PQC qubits, the rate is extremely low, i.e., $\Delta \sim 0.00024 \times \log(M)$. Moreover, we could not observe any critical difference in classification accuracy against the reference, i.e., theoretical accuracy obtained by convex optimization. Therefore, at least for $M \leq 2^{13}$, we conclude that run-time complexity of $\mathcal{O}(M \text{polylog}(M))$ with bounded

classification accuracy is empirically achievable for VQASVM due to $\mathcal{O}(\text{polylog}(M))$ number of parameters (see Supplementary Fig. S11 online for a visual summarization of the result).

Discussion

In this work, we propose a novel quantum-classical hybrid supervised machine learning algorithm that achieves run-time complexity of $\mathcal{O}(M \text{polylog}(M))$, whereas the complexity of the modern classical algorithm is $\mathcal{O}(M^2)$. The main idea of our VQASVM algorithm is to encode optimization parameters that represent the normalized weight for each training data instance in a quantum state using exponentially fewer parameters. We numerically confirmed the convergence and feasibility of VQASVM even in the presence of expectation value estimation error using SPSA. We also observed the sub-quadratic asymptotic run-time complexity of VQASVM numerically. Finally, VQASVM was tested on cloud-based NISQ processors with a toy example dataset to highlight its practical application potential.

Based on the numerical results, we presume that our variational algorithm can bypass the issues evoked by the expressibility³⁶ and trainability relationship of PQCs; i.e., the highly expressive PQC is not likely to be trained with VQA due to the vanishing gradient variance⁴⁵, a phenomenon known as the barren plateau⁴⁶. This problem has become a critical barrier for most VQAs utilizing a PQC to generate solution states. However, given that the SVM is learnable (i.e., the singular values of the kernel matrix decay exponentially⁵), only a few Lagrange multipliers corresponding to the significant support vectors are critically large; $\alpha_i \gg 1/M^{29,30}$. For example, Fig. 4d illustrates the statement. Thus, the PQCs encoding optimal multipliers should not necessarily be highly expressive. We speculate that there exists an ansatz generating these sparse probability distributions. Moreover, the optimal solution would have exponential degeneracy because only measurement probability matters instead of the amplitude of the state itself. Therefore, we could not observe a critical decrease in classification accuracy for these reasons even though the barren plateau exists, i.e., $\Delta = \mathcal{O}(\log(M))$. More analytic discussion on the trainability of VQASVM and the ansatz generating sparse probability distribution should be investigated in the future.

The VQASVM method manifests distinctive features compared to LS-QSVM, which solves the linear algebraic optimization problem of least-square SVM⁴⁷ with the quantum algorithm for linear systems of equations (HHL)²⁰. Given that the fault-tolerant universal quantum computers and efficient quantum data loading with QRAM^{2,22} are possible, the run-time complexity of LS-QSVM is exponentially low: $\mathcal{O}(\kappa_{\text{eff}}^3 \varepsilon^{-3} \log(MN))$ with error ε and effective condition number κ_{eff} . However, near-term implementation of LS-QSVM is infeasible due to lengthy quantum subroutines, which VQASVM has managed to avoid. Also, training LS-QSVM has to be repeated for each query of unseen data because the solution state collapses after the measurements at the end; transferring the solution state to classify multiple test data violates the no-cloning theorem. VQASVM can overcome these drawbacks. VQASVM is composed of much shorter operations; VQASVM circuits are much shallower than HHL circuits with the same moderate system size when decomposed in the same universal gate set. The classification phase of VQASVM can be separated from the training phase and performed simultaneously; training results are classically saved and transferred to a decision circuit in other quantum processing units (QPUs).

We continue the discussion on the advantage of our method compared to other quantum kernel-based algorithms, such as a variational quantum classifier (VQC) and quantum kernel estimator (QKE), which are expected to be realized in the near-term NISQ devices⁸. VQC estimates the label of data \mathbf{x} as $\hat{y} = \text{sgn}\{\tilde{f}(\mathbf{x}; \boldsymbol{\theta}) + b\}$, where $\tilde{f}(\mathbf{x}; \boldsymbol{\theta})$ is the empirical average of a binary function $f(z)$ such that z is the N -bit computational basis measurement result of quantum circuit $W(\boldsymbol{\theta})|\phi(\mathbf{x})\rangle$. The parameters $\boldsymbol{\theta}$ and b are trained with variational methods that minimize the empirical risk: $R_{\text{emp}}(\boldsymbol{\theta}, b) = \sum_{(\mathbf{x}, y) \in \mathcal{S}} \Pr[y \neq \text{sgn}\{\tilde{f}(\mathbf{x}; \boldsymbol{\theta}) + b\}] / |\mathcal{S}|$. This requires $\mathcal{O}(MN \times |\boldsymbol{\theta}|)$ quantum circuit measurements per iteration⁴. Subsequently, the complexity of VQC would match VQASVM from the heuristic point of view. However, VQC does not take advantage of the strong duality; the optimal state that $W(\boldsymbol{\theta}^*)^\dagger$ should generate has no characteristic, whereas the optimal distribution of α_i^* that ansatz of VQASVM should generate is very sparse, i.e., most α_i s are close to zero. Therefore, optimizing VQC in terms of the sufficiently large number of qubits would be vulnerable to issues such as local minima and barren plateaus. On the other hand, QKE estimates the kernel matrix elements \hat{K}_{ij} from the empirical probability of measuring a N -bit zero sequence on quantum circuit $U_{\phi(\mathbf{x}_i)}^\dagger U_{\phi(\mathbf{x}_j)}|\mathbf{0}\rangle$, given the kernel matrix $K_{ij} = |\langle \phi(\mathbf{x}_i) | \phi(\mathbf{x}_j) \rangle|^2$. The estimated kernel matrix is feed into the classical kernel-based algorithms, such as SVM. That is why we used QKE as the reference for the numerical analysis. The kernel matrix can be estimated in $\mathcal{O}(\varepsilon^{-2}M^4)$ measurements with the bounded error of $\|K - \hat{K}\| \leq \varepsilon$. Thus, QKE has much higher complexity than both VQASVM and classical SVM⁸. In addition, unlike QKE, the generalized error converges to zero as $M \rightarrow \infty$ due to the exponentially fewer parameters of VQASVM, strengthening the reliability of the training⁴⁸. The numerical linearity relation with the decision function error \mathcal{E}_f and Δ supports the claim (see Supplementary Fig. S11 online).

VQASVM can be enhanced further with kernel optimization. Like other quantum-kernel-based methods, the choice of the quantum feature map is crucial for VQASVM. Unlike previous methods (e.g., quantum kernel alignment⁴⁹), VQASVM can optimize a quantum feature map online during the training process. Given that $U_{\phi(\cdot)}$ is tuned with other parameters $\boldsymbol{\varphi}$, optimal parameters should be the saddle point: $(\boldsymbol{\theta}^*, \boldsymbol{\varphi}^*) = \arg \min_{\boldsymbol{\theta}} \max_{\boldsymbol{\varphi}} \mathcal{L}_{\phi[\boldsymbol{\varphi}], \lambda}(\boldsymbol{\theta}) + C^{-1} \mathcal{R}(\boldsymbol{\theta})$. In addition, tailored quantum kernels (e.g., $k(\cdot, \cdot) = |\langle \phi(\cdot) | \phi(\cdot) \rangle|^{2r}$) can be adapted with the simple modification¹⁰ on the quantum circuits for VQASVM to improve classification accuracy. However, because the quantum advantage in classification accuracy derived from the power of quantum kernels is not the scope of this paper, we leave the remaining discussion for the future. Another method to improve VQASVM is boosting. Since VQASVM is not a convex problem, the performance may depend on the initial point and not be immune to the overfitting problem, like other kernel-based algorithms. A boosting method can be applied to improve classification accuracy by cascading low-performance

VQASVMs. Because each VQASVM model only requires $\mathcal{O}(\log(M) + N)$ space, ensemble methods such as boosting are suitable for VQASVM.

Methods

Proof of Eq. (13). First, we note that a SWAP test operation ($H_a \cdot \text{cSWAP}_{a \rightarrow b, c} \cdot H_a$) in Fig. 1 measures the Hilbert–Schmidt inner product between two pure states by estimating $\langle \sigma_z^a \rangle = |\langle \phi | \psi \rangle|^2$, where a is the control qubit and $|\phi\rangle$ and $|\psi\rangle$ are states on target qubits b and c , respectively. Quantum registers i and j in Fig. 1 are traced out because measurements are performed on only a and y qubits. The reduced density matrix on x and y quantum registers before the controlled-SWAP operation is $\rho_{x,y} = \sum_{i=0}^{M-1} \alpha_i |\phi(\mathbf{x}_i)\rangle \langle \phi(\mathbf{x}_i)| \otimes |y_i\rangle \langle y_i|$, which is the statistical sum of quantum states $|\phi(\mathbf{x}_i)\rangle \otimes |y_i\rangle$ with probability α_i . Let us first consider the decision circuit (Fig. 1b). Given that the states $|\phi(\mathbf{x}_i)\rangle_x \otimes |y_i\rangle_y$ and $|\phi(\hat{\mathbf{x}})\rangle_{\hat{x}}$ are prepared, $\langle \mathcal{A}_{a \rightarrow x, \hat{x}} \sigma_z^y \rangle = \langle \mathcal{A}_{a \rightarrow x, \hat{x}} \rangle \langle \sigma_z^y \rangle = y_i \langle \mathcal{A}_{a \rightarrow x, \hat{x}} \rangle$ due to separability. Here, $\mathcal{A}_{a \rightarrow x, \hat{x}}$ can be $(H_a \cdot \text{cSWAP}_{a \rightarrow x, \hat{x}} \cdot H_a)^\dagger \sigma_z^a (H_a \cdot \text{cSWAP}_{a \rightarrow x, \hat{x}} \cdot H_a)$ or $\frac{1}{2} \sigma_a^0$. Similarly, for the loss circuit (Fig. 1a), we have states $|\phi(\mathbf{x}_i)\rangle_{x_0} \otimes |y_i\rangle_{y_0}$ and $|\phi(\mathbf{x}_j)\rangle_{x_1} \otimes |y_j\rangle_{y_1}$ with probability $\alpha_i \alpha_j$ such that $\langle \mathcal{A}_{a \rightarrow x_0, x_1} \sigma_z^{y_0} \sigma_z^{y_1} \rangle = \langle \mathcal{A}_{a \rightarrow x_0, x_1} \rangle \langle \sigma_z^{y_0} \sigma_z^{y_1} \rangle = y_i y_j \langle \mathcal{A}_{a \rightarrow x_0, x_1} \rangle$. Therefore, from the definition of our quantum kernel, each term in Eq. (13) matches the loss and decision functions in Eqs. (11) and (12). More direct proof is provided in Ref.^{10,11} and in the Supplementary Information section B (online).

Realization of quantum circuits. In this article, $\mathcal{U}_{\phi, S}$ is realized using uniformly controlled one-qubit gates, which require at most $M - 1$ CNOT gates, M one-qubit gates, and a single diagonal $(\log(M) + 1)$ -qubit gate^{35,50}. We compiled the quantum feature map with a basis gate set composed of Pauli rotations and CNOT. $\mathcal{U}_{\phi, S}$ can be efficiently implemented by replacing all Pauli rotations with uniformly controlled Pauli rotations. The training data label embedding of $\mathcal{U}_{\phi, S}$ can also be easily implemented using a uniformly controlled Pauli X rotation (i.e., setting the rotation angle to π if the label is positive and 0 otherwise). Although this procedure allows one to incorporate existing quantum feature maps, the complexity can increase to $\mathcal{O}(MN^2)$ if the quantum feature map contains all-to-all connecting parameterized two-qubit gates. Nonetheless, such a value of $\mathcal{U}_{\phi, S}$ has linear complexity proportional to the number of training data instances.

Application to IBM quantum processors. Because IBM quantum processors are based on superconducting qubits, all-to-all connections are not possible. Additional SWAP operations among qubits for distant interactions would shorten the effective decoherence time and increase the noise. We carefully selected the physical qubits of *ibmq_montreal* in order to reduce the number of SWAP operations. For $M = 4$ and single qubit embedding, $m = 2$ and $n = 1$. Thus, multi-qubit interaction is required for the following connections: (a, x_0, x_1) , (i_0, i_1, x_0) , (j_0, j_1, x_1) , (i_0, i_1, y_0) , and (j_0, j_1, y_1) . We initially selected nine qubits connected in a linear topology such that the overall estimated single and two-qubit gate errors are lowest among all other possible options. The noise parameters and topology of *ibmq_montreal* are provided by IBM Quantum. For instance, the physical qubits indexed as 1, 2, 3, 4, 5, 8, 11, 14, and 16 in *ibmq_montreal* were selected in this article (see Supplementary Fig. S1 online) We then assign a virtual qubit in the order of $y_0, i_0, i_1, x_0, a, x_1, j_0, j_1, y_1$ so that the aforementioned required connections can be made between qubits next to each other. In conclusion, mapping from virtual qubits to physical qubits proceeds as $\{a \mapsto 5, i_0 \mapsto 2, i_1 \mapsto 1, x_i \mapsto 3, y_i \mapsto 4, j_0 \mapsto 11, j_1 \mapsto 14, x_j \mapsto 8, y_j \mapsto 16\}$ in this experiment. We report that with this arrangement, the circuit depths of loss and decision circuits are correspondingly 60 and 59 for a balanced dataset and 64 and 63 for the an unbalanced dataset in the basis gate set of *ibmq_montreal*: $\{R_z, \sqrt{X}, X, \text{CNOT}\}$.

Additional techniques on SPSSA. The conventional SPSSA algorithm has been adjusted for faster and better convergence. First, the *blocking* technique was introduced. Assuming that the variance σ^2 of objective function $\mathcal{L}_{\phi, \lambda} + C^{-1} \mathcal{R}$ is uniform on parameter θ , the next iteration $t + 1$ is rejected if $[\mathcal{L}_{\phi, \lambda} + C^{-1} \mathcal{R}](\theta^{t+1}) \geq [\mathcal{L}_{\phi, \lambda} + C^{-1} \mathcal{R}](\theta^t) + 2\sigma$. SPSSA would converge more rapidly with blocking by preventing its objective function from becoming too large with some probability (see Supplementary Fig. S10). Second, *Early-stopping* is applied. Iterations are terminated if certain conditions are satisfied. Specifically, we stop SPSSA optimization if the average of last 16 recorded training loss values is greater than or equal to the last 32 recorded values. Early stopping reduces the training time drastically, especially when running on a QPU. Last, we averaged the last 16 recorded parameters to yield the result $\theta^* = \frac{1}{16} \sum_{i=0}^{15} \theta^{t-i}$. Combinations of these techniques were selected for better optimization. We adopted all these techniques for the experiments and simulations as the default condition.

Warm-start optimization. We report cases in which the optimization of IBM Q Quantum Processors yields vanishing kernel amplitudes due to the constantly varying error map problem. The total run time should be minimized to avoid this problem. Because accessing a QPU takes a relatively long queue time, we apply a ‘warm-start’ technique, which reduces number of QPU uses. First, we initialize and proceed a few iterations (32) with a noisy simulation on a CPU and then evaluate the functions on a QPU for the remaining iterations. Note that an SPSSA optimizer requires heavy initialization computation, such as when the initial variance is calculated. With this warm-start method, we are able to obtain better results on some trials.

Data availability

The numerical data generated in this work are available from the corresponding author upon reasonable request. <https://github.com/Siheon-Park/QUIC-Projects>.

References

- Sharma, K. *et al.* Reformulation of the no-free-lunch theorem for entangled datasets. *Phys. Rev. Lett.* **128**, 070501. <https://doi.org/10.1103/PhysRevLett.128.070501> (2022).
- Lloyd, S., Mohseni, M. & Rebentrost, P. Quantum algorithms for supervised and unsupervised machine learning (arXiv preprint) (2013).
- Biamonte, J. *et al.* Quantum machine learning. *Nature* **549**, 195–202. <https://doi.org/10.1038/nature23474> (2017).
- Schuld, M. & Petruccione, F. *Supervised Learning with Quantum Computers* Vol. 17 (Springer, 2018).
- Rebentrost, P., Mohseni, M. & Lloyd, S. Quantum support vector machine for big data classification. *Phys. Rev. Lett.* **113**, 130503. <https://doi.org/10.1103/PhysRevLett.113.130503> (2014).
- Schuld, M., Fingerhuth, M. & Petruccione, F. Implementing a distance-based classifier with a quantum interference circuit. *EPL (Europhys. Lett.)* **119**, 60002. <https://doi.org/10.1209/0295-5075/119/60002> (2017).
- Schuld, M. & Killoran, N. Quantum machine learning in feature Hilbert spaces. *Phys. Rev. Lett.* **122**, 040504. <https://doi.org/10.1103/PhysRevLett.122.040504> (2019).
- Havlíček, V. *et al.* Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212. <https://doi.org/10.1038/s41586-019-0980-2> (2019).
- Lloyd, S., Schuld, M., Ijaz, A., Izaac, J. & Killoran, N. Quantum embeddings for machine learning (arXiv preprint) (2020).
- Blank, C., Park, D. K., Rhee, J.-K.K. & Petruccione, F. Quantum classifier with tailored quantum kernel. *NPJ Quantum Inf.* **6**, 41. <https://doi.org/10.1038/s41534-020-0272-6> (2020).
- Park, D. K., Blank, C. & Petruccione, F. The theory of the quantum kernel-based binary classifier. *Phys. Lett. A* **384**, 126422. <https://doi.org/10.1016/j.physleta.2020.126422> (2020).
- Schuld, M. Supervised quantum machine learning models are kernel methods (arXiv preprint) (2021).
- Liu, Y., Arunachalam, S. & Temme, K. A rigorous and robust quantum speed-up in supervised machine learning. *Nat. Phys.* **17**, 1013–1017. <https://doi.org/10.1038/s41567-021-01287-z> (2021).
- Blank, C., da Silva, A. J., de Albuquerque, L. P., Petruccione, F. & Park, D. K. Compact quantum kernel-based binary classifier. *Quantum Sci. Technol.* **7**, 045007 (2022).
- Cortes, C. & Vapnik, V. Support-vector networks. *Mach. Learn.* **20**, 273–297. <https://doi.org/10.1007/BF00994018> (1995).
- Boser, B. E., Guyon, I. M. & Vapnik, V. N. A training algorithm for optimal margin classifiers. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, 144–152 (1992).
- Vapnik, V., Golowich, S. E. & Smola, A. Support vector method for function approximation, regression estimation, and signal processing. *Adv. Neural Inf. Process. Syst.* **20**, 281–287 (1997).
- Guyon, I., Vapnik, V., Boser, B., Bottou, L. & Solla, S. Capacity control in linear classifiers for pattern recognition. In *Proceedings of 11th IAPR International Conference on Pattern Recognition. Vol. II. Conference B: Pattern Recognition Methodology and Systems*, 385–388. <https://doi.org/10.1109/ICPR.1992.201798> (IEEE Comput. Soc. Press, 1992).
- Lloyd, S., Mohseni, M. & Rebentrost, P. Quantum principal component analysis. *Nat. Phys.* **10**, 631–633. <https://doi.org/10.1038/nphys3029> (2014).
- Harrow, A. W., Hassidim, A. & Lloyd, S. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.* **103**, 150502. <https://doi.org/10.1103/PhysRevLett.103.150502> (2009).
- Aaronson, S. Read the fine print. *Nat. Phys.* **11**, 291–293. <https://doi.org/10.1038/nphys3272> (2015).
- Giovannetti, V., Lloyd, S. & Maccone, L. Quantum random access memory. *Phys. Rev. Lett.* **100**, 160501. <https://doi.org/10.1103/PhysRevLett.100.160501> (2008).
- Preskill, J. Quantum computing in the NISQ era and beyond. *Quantum* **2**, 79. <https://doi.org/10.22331/q-2018-08-06-79> (2018).
- Benedetti, M., Lloyd, E., Sack, S. & Fiorentini, M. Parameterized quantum circuits as machine learning models. *Quantum Sci. Technol.* **4**, 043001 (2019).
- Cerezo, M. *et al.* Variational quantum algorithms. *Nat. Rev. Phys.* **3**, 625–644. <https://doi.org/10.1038/s42254-021-00348-9> (2021).
- Chang, C.-C. & Lin, C.-J. LIBSVM. *ACM Trans. Intell. Syst. Technol.* **2**, 1–27. <https://doi.org/10.1145/1961189.1961199> (2011).
- Coppersmith, D. & Winograd, S. Matrix multiplication via arithmetic progressions. *J. Symb. Comput.* **9**, 251–280. [https://doi.org/10.1016/S0747-7171\(08\)80013-2](https://doi.org/10.1016/S0747-7171(08)80013-2) (1990).
- Giuntini, R. *et al.* Quantum state discrimination for supervised classification (arXiv preprint) (2021).
- Boyd, S. P. & Vandenberghe, L. *Convex Optimization* (Cambridge University Press, 2004).
- Deisenroth, M. P., Faisal, A. A. & Ong, C. S. *Mathematics for Machine Learning* (Cambridge University Press, 2020).
- Mangasarian, O. & Musicant, D. Successive overrelaxation for support vector machines. *IEEE Trans. Neural Netw.* **10**, 1032–1037. <https://doi.org/10.1109/72.788643> (1999).
- Frie, T.-T., Cristianini, N. & Campbell, C. The kernel-adatron algorithm: A fast and simple learning procedure for support vector machines. In *Machine Learning: Proceedings of the Fifteenth International Conference (ICML'98)*, 188–196 (1998).
- Hsu, C.-W. & Lin, C.-J. A simple decomposition method for support vector machines. *Mach. Learn.* **46**, 291–314. <https://doi.org/10.1023/A:1012427100071> (2002).
- Hsu, Chih-Wei. & Lin, Chih-Jen. A comparison of methods for multiclass support vector machines. *IEEE Trans. Neural Netw.* **13**, 415–425. <https://doi.org/10.1109/72.991427> (2002).
- Mottonen, M., Vartiainen, J. J., Bergholm, V. & Salomaa, M. M. Transformation of quantum states using uniformly controlled rotations (arXiv preprint) (2004).
- Sim, S., Johnson, P. D. & Aspuru-Guzik, A. Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms. *Adv. Quantum Technol.* **2**, 1900070. <https://doi.org/10.1002/qute.201900070> (2019).
- Kandala, A. *et al.* Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets. *Nature* **549**, 242–246. <https://doi.org/10.1038/nature23879> (2017).
- Mitarai, K., Negoro, M., Kitagawa, M. & Fujii, K. Quantum circuit learning. *Phys. Rev. A* **98**, 032309. <https://doi.org/10.1103/PhysRevA.98.032309> (2018).
- Schuld, M., Bergholm, V., Gogolin, C., Izaac, J. & Killoran, N. Evaluating analytic gradients on quantum hardware. *Phys. Rev. A* **99**, 032331. <https://doi.org/10.1103/PhysRevA.99.032331> (2019).
- ANIS, M. D. S. *et al.* Qiskit: An open-source framework for quantum computing. <https://doi.org/10.5281/zenodo.2573505> (2021).
- Spall, J. C. A one-measurement form of simultaneous perturbation stochastic approximation. *Automatica* **33**, 109–112 (1997).
- Spall, J. Adaptive stochastic approximation by the simultaneous perturbation method. *IEEE Trans. Autom. Control* **45**, 1839–1853. <https://doi.org/10.1109/TAC.2000.880982> (2000).
- Fisher, R. A. & Marshall, M. Iris data set (1936).
- Deng, L. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Process. Mag.* **29**, 141–142 (2012).
- Holmes, Z., Sharma, K., Cerezo, M. & Coles, P. J. Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum* **3**, 010313. <https://doi.org/10.1103/PRXQuantum.3.010313> (2022).

46. McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R. & Neven, H. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.* **9**, 4812. <https://doi.org/10.1038/s41467-018-07090-4> (2018).
47. Suykens, J. A. K. & Vandewalle, J. Least squares support vector machine classifiers. *Neural Process. Lett.* **9**, 293–300. <https://doi.org/10.1023/A:1018628609742> (1999).
48. Caro, M. C. *et al.* Generalization in quantum machine learning from few training data. *Nat. Commun.* **13**, 1–11 (2022).
49. Glick, J. R. *et al.* Covariant quantum kernels for data with group structure (arXiv preprint). <https://doi.org/10.48550/arxiv.2105.03406> (2021).
50. Bergholm, V., Vartiainen, J. J., Möttönen, M. & Salomaa, M. M. Quantum circuits with uniformly controlled one-qubit gates. *Phys. Rev. A* **71**, 052330. <https://doi.org/10.1103/PhysRevA.71.052330> (2005).

Acknowledgements

This work was supported by the Samsung Research Funding & Incubation Center of Samsung Electronics under Project Number SRFC-TF2003-01. We acknowledge the use of IBM Quantum services for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Quantum team.

Author contributions

S.P. contributed to the development and experimental verification of the theoretical and circuit models; D.K.P. contributed to developing the initial concept and the experimental verification; J.K.R. contributed to the development and validation of the main concept and organization of this work. All co-authors contributed to the writing of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-023-29495-y>.

Correspondence and requests for materials should be addressed to J.-K.K.R.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023