Importance of Kernel Bandwidth in Quantum Machine Learning

Ruslan Shaydulin (✉ ruslan@anl.gov)

Argonne National Laboratory

Stefan Wild

Argonne National Laboratory

Research Article

Keywords: Quantum Machine Learning, Kernel Bandwidth, Importance, Algorithms, kernel methods

Posted Date: December 3rd, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-1087188/v1>

License: This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Importance of Kernel Bandwidth in Quantum Machine Learning

Ruslan Shaydulin^{1,*} and Stefan M. Wild¹

¹Mathematics and Computer Science Division, Argonne National Laboratory, Lemont, IL 60439 USA

*ruslan@anl.gov

ABSTRACT

Quantum kernel methods are considered a promising avenue for applying quantum computers to machine learning problems. However, recent results overlook the central role hyperparameters play in determining the performance of machine learning methods. In this work we identify the hyperparameter controlling the bandwidth of a quantum kernel and show that it controls the expressivity of the resulting model. We use extensive numerical experiments with multiple quantum kernels and classical datasets to show consistent change in the model behavior from underfitting (bandwidth too large) to overfitting (bandwidth too small), with optimal generalization in between. We draw a connection between the bandwidth of classical and quantum kernels and show analogous behavior in both cases. Furthermore, we show that optimizing the bandwidth can help mitigate the exponential decay of kernel values with qubit count, which is the cause behind recent observations that the performance of quantum kernel methods decreases with qubit count. We reproduce these negative results and show that if the kernel bandwidth is optimized, the performance instead improves with growing qubit count and becomes competitive with the best classical methods.

Introduction

Algorithms designed for quantum computers are theoretically able to provide exponential speedups over the best known classical algorithms for certain problems, most famously integer factoring¹. Recent advances in quantum computing hardware open the possibility of realizing this potential. Urgently needed, however, is development of novel algorithms that effectively leverage the power of quantum computation. A particularly promising problem domain is machine learning, owing to the ubiquity of machine learning problems in science and technology.

A number of approaches have been proposed for applying quantum computers to machine learning problems. In this work we focus on a subset of such approaches that can be reformulated as kernel methods. The central idea of kernel methods is to embed the data into a feature space (typically, of dimension higher than the input data) in which it becomes easier to analyze. The analysis is then performed by using solely the values of similarities (kernel values) between the representations of the data points in feature space. Most supervised machine learning methods (quantum or classical) can be equivalently reformulated as kernel methods with an appropriately defined (potentially, quantum) kernel^{2,3}. Quantum kernel methods^{4,5} have been shown to theoretically provide speedups over classical methods^{6,7}. In quantum kernel methods, datapoints are mapped to quantum states using a quantum feature map, and the value of the kernel between two datapoints is given by some similarity measure (such as fidelity) of the corresponding quantum states. The power of quantum kernel methods comes from being able to process the data using an exponentially sized Hilbert space, performing computations that are hard classically even for sizes attainable on currently available quantum hardware⁸.

Recently, the exponential dimensionality of the space into which the classical data is being mapped has been identified as a potential obstacle to achieving quantum advantage using kernel methods. The fidelity of two random quantum states decreases exponentially with the number of qubits (typically set to be equal to the number of dimensions of the input data), leading to the exponential vanishing of kernel values and making learning impossible. This behavior is analogous to the curse of dimensionality in classical kernel methods⁹ and not specific to the choice of fidelity as the similarity measure. One recently proposed approach to overcome this limitation is controlling the inductive bias of the quantum kernel methods by projecting the quantum state into a lower-dimensional subspace^{7,10}. In general, however, additional information is required to appropriately choose the projection¹⁰.

Prior works consider the quantum feature maps as fixed. This approach, however, overlooks the important role hyperparameters play in classical machine learning in general and in classical kernel methods in particular³, and the potential for hyperparameters to control the inductive bias of the method and overcome the limitations stemming from exponential dimensionality of the space. One such hyperparameter is kernel bandwidth, which is known to affect the performance of methods such as support vector machines (SVMs) and is routinely optimized when SVMs are used in practice.

In this work we identify quantum kernel bandwidth as a centrally important hyperparameter for quantum kernel methods. We demonstrate that varying it controls the expressiveness of the model, with high bandwidth leading to underfitting and low bandwidth to overfitting. We show that for the quantum feature maps considered, the bandwidth can be controlled by rescaling the datapoints. We draw an analogy with the bandwidth of classical radial basis function (RBF) kernels and show that varying bandwidth affects expressiveness of the model in a similar way in both quantum and classical cases. As a consequence of the effect of the bandwidth on generalization, optimizing the bandwidth can improve the performance from being comparable to a random guess to being competitive with the best classical methods. We reproduce recent numerical results used to support an argument against quantum kernels⁷. We show how changing the value of the kernel bandwidth can lead to significantly better performance than the results in⁷. Moreover, we observe that when the bandwidth is optimized, the performance of quantum kernel methods improves with qubit count, which is the opposite scaling behavior from the one documented in⁷. We provide extensive numerical evidence of the importance of the bandwidth and its effect on model generalization and performance with varying qubit count, performing simulations with multiple quantum feature maps and datasets using up to 26 qubits. We show that hardware limitations, such as finite precision of controls and the variance introduced by sampling, are not obstacles to achieving the observed performance.

Results

Setup

We begin by formalizing our notions of quantum kernels, quantum feature maps, and machine learning problems, as well as establishing notation. We consider the problem of supervised learning, specifically the task of classification, although we expect our conclusions to apply more generally. Given a training dataset of N pairs $\{(\mathbf{x}_i, y_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathbb{R}^d$ is a datapoint and $y_i \in \{0, 1\}$ is a binary label produced by some unknown map $m : \mathbb{R}^d \rightarrow \{0, 1\}$, the goal is to learn a map from datapoints to labels that agrees with the true map m with high probability on an unseen test set.

We consider quantum models where a datapoint $\mathbf{x}_i$ is encoded in a quantum state $|\mathbf{x}_i\rangle$ by a parameterized unitary: $|\mathbf{x}_i\rangle = U_{\text{enc}}(\mathbf{x}_i)|0\rangle$. We will refer to the unitary U_{enc} as a quantum feature map. A kernel matrix K is obtained by computing $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j) = |\langle \mathbf{x}_i | \mathbf{x}_j \rangle|^2$ (the quantum kernel) for all pairs of datapoints. This value can be computed on a quantum computer by measuring the value of observable $|0\rangle\langle 0|$ on the state $U_{\text{enc}}^\dagger(\mathbf{x}_i)U_{\text{enc}}(\mathbf{x}_j)|0\rangle$. This kernel matrix is then used inside an SVM or other kernel methods^{11,12}. A quantum feature map and a classical kernel method such as an SVM fully define a quantum kernel method.

In this paper we use the quantum kernel matrix in a support vector classifier (SVC). SVC seeks an optimal separating hyperplane between two (potentially nonseparable) classes in the feature space. An optimal hyperplane is one that maximizes the distance (margin) to the closest point from either class¹³. The hyperplane is found by determining coefficients $\boldsymbol{\alpha} \in \mathbb{R}^N$ that maximize the objective function

$$\sum_{i=1}^N \alpha_i - \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \alpha_i \alpha_j y_i y_j k(\mathbf{x}_i, \mathbf{x}_j), \quad (1)$$

subject to $\sum_{i=1}^N \alpha_i y_i = 0$ and $0 \leq \alpha_i \leq C$, $i = 1, \dots, N$. The penalty term $C > 0$ controls the trade-off between the goals of minimizing the training error and maximizing the margin: smaller values of C lead to larger margin at the cost of higher training error.

Motivating example

The recent results on quantum kernel methods do not explicitly study the choice of hyperparameters in quantum feature maps. As a motivating example, we consider the paper⁷, which includes numerical experiments on the performance of the quantum kernel method with the feature map given by

$$|\mathbf{x}_i\rangle = \left(\prod_{j=1}^d \exp \left(-i \frac{t}{T} x_{ij} (X_j X_{j+1} + Y_j Y_{j+1} + Z_j Z_{j+1}) \right) \right)^{T_{n+1}} \bigotimes_{j=1}^{n+1} |\psi_j\rangle. \quad (2)$$

Here X_j , Y_j , and Z_j are the Pauli operators acting on qubit j , and $|\psi_j\rangle$ is a Haar-random single-qubit state. The Haar-random state $|\psi_j\rangle$ is sampled and fixed for each qubit. This feature map represents a d -dimensional datapoint as a $d+1$ -qubit quantum state¹⁴. We refer to this feature map as the Hamiltonian evolution feature map.

The feature map given by equation 2 has two hyperparameters: the total evolution time t (equivalent to the scaling of the inputs $\mathbf{x}_i \leftarrow t\mathbf{x}_i$) and the number of Trotter steps T . Huang et al.⁷ set $t = \frac{d}{3}$ and $T = 20$ and reported prediction accuracy on the test set of the quantum kernel method (SVC) with this feature map on the Fashion-MNIST (`fmnist`) dataset¹⁵. Each datapoint

Figure 1. (a) Schematic overview of the effect of bandwidth on kernel method performance. Some bandwidth optimization is typically required to avoid under- or overfitting. (b) Prediction accuracy of the quantum kernel method with Hamiltonian evolution feature map as a function of number of qubits. “Reported” and “Reproduced” results have $t = \frac{d}{3}$, where d is the dimensionality of the datapoint. “Optimized” presents performance of the kernel method with the same feature map and $t = 0.05$. The performance of our quantum kernel with optimized bandwidth approaches that of the best classical methods reported in⁷. Note that from the construction given in Ref.⁷ (see equation 2), an n -dimensional datapoint is embedded into an $(n + 1)$ -qubit state.

is reduced to the specified number of dimensions by using principal component analysis (PCA). In Fig. 1b we present the results reported in⁷ and the results we reproduced independently. The details of the SVC implementation and data preprocessing are given in the Methods section. Our reproduced results are presented as boxplots over 20 choices of the Haar-random initial state fixed across all qubit counts, with the box showing the middle 50% quantile range and whiskers showing minimum and maximum. The reported and reproduced results agree well, and both show that the performance of the quantum kernel method decreases with increasing qubit count, approaching the performance of a random guess (i.e., 0.5 accuracy) at about 20 qubits.

While Huang et al.⁷ performed extensive hyperparameter optimization of both the SVC that uses the quantum kernel and the classical methods with which the quantum kernel method is compared, no hyperparameter optimization is reported for the quantum feature map itself. In fact, choosing $t = 0.05$ and keeping everything else fixed lead to much better performance, approaching the performance of the best classical machine learning methods (see Fig. 1b). The reason is that the parameter t controls the bandwidth of the quantum kernel. Large values of t lead to the kernel’s being too “narrow” (small bandwidth), making learning impossible. We discuss the relationship between the scaling factor, quantum kernel bandwidth, and model performance in detail below.

In addition to providing improved performance, this hyperparameter choice leads to an opposing conclusion about the scaling of the performance with qubit count: in this case, the performance increases with qubit count until plateauing at about 14 qubits. The improvement in performance can be partially explained by observing that with larger qubit count, more components are kept in the PCA-based reduction, thereby giving the model access to a larger part of the explanation of variance in the data. Unlike the previous results^{7,10}, we do not observe a large drop in performance for larger qubit counts. This example highlights the importance of hyperparameter selection in determining the effectiveness of quantum kernel methods.

Importance of hyperparameters in quantum machine learning

We now show that the ability of hyperparameter choice to make or break performance of a quantum kernel method is not limited to the above example. To this end, we perform a systematic study of how bandwidth affects the prediction accuracy of quantum kernel methods. In addition to the Hamiltonian evolution feature map given in equation 2, we consider a feature map inspired by the instantaneous quantum polynomial (IQP) circuits used in^{5,7}:

$$|\mathbf{x}_i\rangle = U_Z(\mathbf{x}_i) H^{\otimes d} U_Z(\mathbf{x}_i) H^{\otimes d} |0^d\rangle, \quad (3)$$

where

$$U_Z(\mathbf{x}_i) = \exp \left(\sum_{j=1}^d \lambda x_{ij} Z_j + \sum_{j=1}^d \sum_{j'=1}^d \lambda^2 x_{ij} x_{ij'} Z_j Z_{j'} \right). \quad (4)$$

Figure 2. Prediction accuracy of the quantum kernel method with IQP (a) and Hamiltonian evolution (c) feature maps and as a function of the scaling factor. (b) Effects of finite precision of controls on performance of the quantum kernel method. Error bars show one standard deviation. (d) Effects of finite sampling on performance of the quantum kernel method. Top: median value and standard deviation of the non-diagonal entries of the kernel matrix K computed from full-precision quantum state. Bottom: test accuracy for kernel matrix computed from full-precision quantum state and from 5,000 samples as a function of scaling factor applied to the data.

We endow this feature map with a hyperparameter $\lambda > 0$ analogous to the total evolution time in the Hamiltonian evolution feature map. Since both hyperparameters are equivalent to a rescaling of the input data, we refer to them as “scaling factors.” The scaling factor λ can be subsumed into the definition of the input data by setting $\mathbf{x}_i \leftarrow \lambda \mathbf{x}_i$. The scaling factor can be understood as a hyperparameter of the quantum feature map since it can always be included in the definition of the feature map if an assumption on the input data distribution is made (e.g., the data has mean zero and standard deviation one). We show below that the scaling factor behaves analogously to bandwidth in classical kernel methods. The performance is evaluated on three real-world datasets. In addition to the Fashion-MNIST (fmnist) dataset, we use Kuzushiji-MNIST (kmnist), a dataset containing cursive Japanese characters¹⁶, and a cosmological dataset focused on the task of supernova classification (plasticc)¹⁷. Two of the datasets were used to study quantum kernel methods (fmnist in⁷ and plasticc in¹⁸). The datasets are balanced; that is, the number of datapoints with label 0 is approximately equal to the number of datapoints with label 1 for both training and testing datasets. All input data is normalized to be centered around zero with the standard deviation of approximately one. A precise description of the experimental setup is given in the Methods section.

We observe that hyperparameter selection significantly affects the performance of quantum kernel methods. We find in particular that the performance of the models varies drastically with the scaling factor (bandwidth). This relationship is presented in Figs. 2a and 2c. Since no significant effect is observed from varying the number of Trotter steps T in the Hamiltonian evolution feature map, here and for the remainder of the paper we set $T = 40$; results with other values are available in the Supplementary Information. We observe a clear trend in optimal values of the scaling factor, with larger values leading to overfitting and smaller values to underfitting. This relationship is schematically presented in Fig. 1a. Counterintuitively, the scaling factor of one (corresponding to the data having standard deviation of approximately one) is a bad choice for all the datasets and feature maps considered.

We can understand the importance of the scaling factor by considering the relationship between qubit count, scaling factor, and the values of the kernel matrix. The performance of SVC is sensitive to kernel bandwidth. If the kernel bandwidth is too small (kernel is “narrow”), then the kernel matrix will be close to identity; and if the bandwidth is too large (kernel is “wide”), all kernel matrix entries will be close to one. In both scenarios learning is impossible, which is why kernel bandwidth is routinely optimized in classical kernel methods. Scaling factor is analogous to kernel bandwidth for quantum kernels. When the scaling factor is large, the kernel matrix approaches identity as the qubit count grows. Correspondingly, when the scaling

Figure 3. Scaling of the kernel with the IQP quantum feature map. **(a)** Prediction accuracy as a function of number of qubits and the scaling factor λ . For each qubit count, the value of the scaling factor corresponding to the highest accuracy is highlighted with a purple square. The trend in scaling factor values corresponding to the highest accuracy is evident. **(b)** Kernel matrix entries as a function of number of qubits and the scaling factor λ . If $\|K - I\|^2$ is small, non-diagonal kernel matrix entries are small (i.e., K approaches the identity). The value $\|K - I\|^2$ is largest when all kernel matrix entries are close to one. **(c)** The scaling of the kernel as a function of the scaling factor λ . The points $\mathbf{x}_i$ and $\mathbf{x}_j$ are taken from the *fnnist* dataset; the number of qubits is 20. The kernel scales exponentially with λ . **(d)** Value of the kernel of two points as a function of the number of qubits. Exponential decay in the value of the kernel entry with the number of qubits is clearly visible.

factor is small, kernel matrix entries all approach one. This relationship is shown in Fig. 3b. Both of those extremes lead to suboptimal performance (see Fig. 3a); the goal of hyperparameter optimization is therefore to find the “Goldilocks” value of the scaling factor. We discuss the analogy with classical kernel bandwidth and the effect on model generalization in detail below by comparing the behavior of quantum kernels with that of the classical RBF kernel. Previously, in¹⁸ rescaling of the data was shown to lead to an improvement in performance of SVC with a different feature map applied to the *plasticc* dataset.

Let us now consider the impact of scaling factor and qubit count on the values of the kernel matrix. First, note that if the data is rescaled by λ , then the operator (4) becomes

$$U_Z(\lambda \mathbf{x}_i) = \exp \left(\lambda \sum_{j=1}^d x_{ij} Z_j + \lambda^2 \sum_{j=1}^d \sum_{j'=1}^d x_{ij} x_{ij'} Z_j Z_{j'} \right) = \left(\exp \left(\sum_{j=1}^d x_{ij} Z_j + \lambda \sum_{j=1}^d \sum_{j'=1}^d x_{ij} x_{ij'} Z_j Z_{j'} \right) \right)^\lambda.$$

Analogous behavior can be easily shown for the Hamiltonian evolution feature map. This highlights the exponential relationship between the scaling factor and the kernel value for the family of kernels considered in this work. Figure 3c numerically elucidates this exponential relationship by plotting $k(\lambda \mathbf{x}_i, \lambda \mathbf{x}_j)$ against $k(\mathbf{x}_i, \mathbf{x}_j)^\lambda$ for two fixed datapoints. Second, note the exponential relationship between the qubit count and the kernel value. As the number of qubits grows, the fidelity of two random states decays exponentially. Correspondingly, the fidelity of the embeddings $|\mathbf{x}_i\rangle$ of the datapoints, equal to the value of the quantum kernel of the datapoints, also decays exponentially. This can be observed in Fig. 3d, which shows the scaling of the value of the kernel of two points with the number of qubits. The points in Fig. 3d are sampled from a normal distribution with mean zero and standard deviation of one (“Normal”) or are taken from the *fnnist* dataset, preprocessed as described in the Methods section, and scaled to have mean zero and standard deviation one. In both cases, the exponential decay in the kernel value with qubit count is clearly visible. These two exponential relationships have important implications since high qubit count is a necessary (but not sufficient) prerequisite for the possibility of quantum advantage. While the kernel value decays exponentially irrespective of the fixed value of the scaling factor, selecting a scaling factor based on the qubit count can mitigate the effects of this decay.

For the range of qubit counts accessible in classical simulation, we find that choosing one scaling factor irrespective of system size leads to good performance. However, we expect that in general the value of scaling factor will vary with qubit count. Initial evidence of this can be glanced from Figs. 3a, wherein the optimized value of the scaling factor decreases slightly for higher qubit count.

As shown in Fig. 1b, the Hamiltonian evolution feature map with optimized scaling factor can achieve accuracy approaching that of the best classical model reported in⁷. The improvement in performance with qubit count, observed in the motivating example, is present for all datasets and qubit counts considered. Figure 4 shows prediction accuracy as a function of datapoint dimension, which is equal to the qubit count in the case of the IQP feature map and is equal to (Number of qubits – 1) for the Hamiltonian evolution feature map. With the scaling factor fixed to the optimized value, we observe improvement in performance with growing qubit count up to a certain threshold at which the performance plateaus. The value of the threshold

Figure 4. Prediction accuracy with varying datapoint dimension for feature maps with optimized hyperparameters for a given dataset. One value of hyperparameters is chosen for each dataset and is fixed across all datapoint dimensions. For the Hamiltonian evolution feature map, mean accuracy over all initial state seeds is presented. The datapoint dimension is equal to the number of qubits for the IQP feature map and is equal to the (number of qubits $- 1$) for the Hamiltonian evolution (“HamEvo”) feature map. For the easier datasets, *fmnist* and *kmnist*, the performance stops improving after a certain threshold is reached, whereas for the harder dataset, *plasticc*, the performance keeps improving up to the largest datapoint dimension considered.

depends on the difficulty of the classification task. For *fmnist*, considered in⁷, the performance stops improving at about 13 qubits, whereas for the more difficult *plasticc* dataset, considered in¹⁸, the performance improves for the entire range of qubit counts considered, namely, up to 22. Similar to the motivating example, the improvement in performance can be partially attributed to the increased dimensionality of the input datapoint, which allows PCA to preserve more data variance. We note, however, that within the regime considered, no deleterious effects of increased dimensionality of the Hilbert space are observed.

Impact of bandwidth on model generalization and comparison with classical kernels

In the preceding section we showed that the performance of quantum kernel methods is sensitive to quantum kernel bandwidth (scaling factor). This sensitivity is due to how the kernel bandwidth affects the expressivity of the model, which we now explain. We begin by noting that the behavior is completely analogous to the variation in the performance of classical kernel methods with kernel bandwidth. The role of kernel bandwidth and its effect on SVC performance are well known in classical machine learning^{11,19,20}. To illustrate its effect, we consider the RBF kernel given by $k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$; here the hyperparameter γ controls the bandwidth. Analogously to the case of the quantum feature map we considered above, γ can be subsumed into the definition of input data as a scaling factor. The kernel matrix $K_{ij} = k(\mathbf{x}_i, \mathbf{x}_j)$ is then used in an SVC analogously to the quantum kernel case described above.

If γ is large, the kernel is “narrow,” and the SVC with such a kernel can fit any labels, which may lead to overfitting. This is analogous to the scaling factor λ above, where large scaling factor values lead to overfitting. Figure 5 shows overfitting resulting from an overly expressive kernel. For large values of γ or λ , the score on the training set goes to one, and the score on the test set decreases to that of random guessing. On the other hand, for sufficiently small values, the kernel is “wide,” and the SVC with such kernel is insufficiently expressive, leading to underfitting. The goal is therefore to identify the bandwidth that makes the model sufficiently but not overly expressive.

To illustrate the effect of quantum kernel bandwidth (controlled by the scaling factor) on the expressiveness of the model, we present two sets of experiments in Fig. 5. We optimize the penalty term C in SVC either using cross-validation or to directly maximize the training score (see Methods for detailed description). When the penalty term that maximizes the training score is chosen, we observe clear under- and overfitting for small and large values of the scaling factor, respectively. Additional information about the performance of the methods without cross-validation is given in the Supplementary Information. In the preceding sections, however, we presented the results with cross-validation since they show slightly better performance. In this case, for the easier *fmnist* and *kmnist* datasets, the behavior is similar. For the harder *plasticc* dataset, the training score for the largest value of the scaling factor λ decreases, thereby indicating the absence of overfitting. We observe that for the *plasticc* dataset, for both feature maps the cross-validation hyperparameter optimization leads to the smallest possible value C , corresponding to the largest margin at the cost of a lower training score. For this particular dataset, this is sufficient to prevent overfitting. For the other datasets considered, the hyperparameter optimization does not lead to the largest penalty, suggesting that increasing the penalty would not help overcome overfitting.

Figure 5. Effects of kernel bandwidth on model generalization on *fmnist* dataset for (a) classical RBF kernel and quantum kernels using 19 qubits with (b) and without (c) cross-validation. Dashed and solid lines represent the training and test scores, respectively. For large values of γ or the scaling factor, the kernel is narrow, and the models can fit any data, leading to overfitting. For small values, the kernel is wide, and the model is insufficiently expressive, leading to underfitting. The maximum average number of support vectors per class is 400 (i.e., every training point is a support vector).

Effects of limited control precision and finite sampling

So far we have demonstrated that with appropriate bandwidth selection, performance of quantum kernel methods need not deteriorate with the growth in qubit count and can approach the performance of the best classical methods. Two major obstacles arise in realizing the observed performance of quantum models in practice: the errors introduced by finite sampling and the limited precision of controls. We now show that these obstacles do not prevent achieving good model performance.

The first obstacle is that, in practice, the kernel matrix is evaluated by performing measurements of a quantum device, limiting the effective precision with which kernel matrix elements can be evaluated. In general, the errors introduced by the finite sampling statistics need not preclude quantum advantage⁶. For small values of the kernel, however, finite sampling may lead to large relative errors²¹ and can potentially nontrivially affect prediction performance. Moreover, the variance of the errors in the value of $|\langle 0|\psi\rangle|^2$ introduced by finite sampling differs depending on the choice of state $|\psi\rangle$, necessitating careful numerical investigation. This can be easily observed by comparing the variance between $|\psi\rangle = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} |x\rangle$ and $|\psi\rangle$ being an n -qubit GHZ state²². To numerically evaluate the effects of finite sampling in a realistic setting, we consider the Hamiltonian evolution feature map with the hyperparameters corresponding to the highest prediction accuracy on the *fmnist* dataset ($n = 19$, $T = 40$). For each value of the scaling factor, we recompute the 5-by-5 subset of the corresponding kernel matrix using 5,000 shots 10 times, and we use the mean of the standard deviations of the submatrix entries as the standard deviation for i.i.d. normal noise with mean zero applied to the full matrix. The noise is applied in a way that preserves the symmetry of the matrix. The resulting matrix may not be positive semi-definite; we take the closest positive semi-definite approximation as the kernel matrix for the SVC.

Figure 2d presents the prediction accuracy obtained with the kernel matrix evaluated from full quantum state up to machine precision (“Exact”) and with sampling noise (“5k shots”). We observe poor performance in the cases where the nondiagonal kernel matrix values are all close to zero or all close to one (i.e., kernel very narrow or very wide). However, an appropriate kernel bandwidth (equivalent to scaling the data) can mitigate this issue by preventing the values of the kernel matrix from concentrating; see Fig. 2d, top.

The second obstacle is finite precision of controls available on noisy intermediate-scale quantum (NISQ) hardware, which limits the range of values that can be used as inputs to the quantum feature map. This can potentially limit the scaling of the data. Note that this issue is limited to NISQ devices, since fault-tolerant architectures allow synthesis of single-qubit operations with precision matching that of classical machines²³. To numerically evaluate the effects of finite precision, we recompute the kernel matrix values with parameters rounded to no more than a fixed number of decimals. Figure 2b shows that with the scaling factor that maximizes performance in the full-precision case (0.05), three decimals of precision is sufficient to match the performance of the noiseless case. As the scaling factor decreases, the number of decimals required to match the full precision

increases. We note that precision of trapped ion hardware currently available on the cloud (three decimals²⁴) is sufficient to saturate the noiseless performance of a scaling factor that is five times smaller than the one maximizing performance at 19 qubits, suggesting that the control precision would not be the limiting factor for the qubit counts expected in the near term.

Discussion

Numerical evaluation of promising quantum algorithms is an essential complement to theoretical analysis. When identifying opportunities and no-go results for quantum advantage, it is therefore crucial to consider all aspects of the quantum algorithm that can be used to boost its performance. In this work we identify quantum kernel bandwidth as the key hyperparameter controlling the expressiveness of the model, with excessively small bandwidth leading to overfitting and large bandwidth leading to underfitting. We show that careful kernel bandwidth selection can mitigate the exponential decay of kernel values with qubit count and avoid overfitting and the corresponding drop-off in prediction accuracy. As a pleasant surprise, we discover that for a given dataset a wide range of bandwidth values work well (this is most clearly visible in Fig. 5). This observation suggests that optimizing kernel bandwidth need not be difficult in practice: finding the correct order of magnitude is often sufficient for achieving good performance.

Recent proposals to control the inductive bias of quantum kernels by projecting them into a lower-dimensional subspace^{7,10} are clearly complementary to the ideas put forward in this work since they attempt to overcome similar challenges. The projection can be straightforwardly combined with bandwidth optimization to more precisely modulate the inductive bias of the model. More interestingly, the projection operation can itself be viewed as a hyperparameter in the absence of sufficient prior knowledge. Such techniques may be required to achieve good performance with qubit counts significantly larger than accessible for study at the time of writing.

In this paper we consider only two hyperparameters—scaling factor and the number of Trotter steps in the Hamiltonian evolution feature map. We see a significant effect for only one of them, the scaling factor, which controls the kernel bandwidth. We anticipate that more elaborate feature maps will be explored with potentially more parameters that can be tuned. Doing so may help bridge the gap between fully trainable quantum embeddings²⁵, which require expensive training on a quantum computer, and fixed feature maps. Careful consideration of hyperparameters of quantum kernels may enable the tuning of the inductive bias of the model in a way that is infeasible classically, providing a potential path to quantum advantage in machine learning.

Methods

Data preprocessing We preprocess the input data in the following way. For `fmnist`¹⁵, we follow the preprocessing pipeline of⁷, which results in the two classes of dresses (class 0) and shirts (class 3). For `kmnist`¹⁶, we arbitrarily choose class 1 and class 4. For `plasticc`¹⁷, we use the dataset preprocessed as described in¹⁸, provided by the authors. For all datasets, each datapoint is rescaled such that the mean value is zero and the standard deviation is one. Then, each datapoint (28-by-28 pixel images for `fmnist` and `kmnist` and 67-dimensional vector for `plasticc`) is reduced to the required number of dimensions by using the scikit-learn²⁶ implementation of PCA. We use 800 randomly selected points as the training set and 200 randomly selected points as the test set. The `fmnist` and `kmnist` datasets provide separate training and test sets, and for `plasticc` we split the dataset into training and testing parts randomly, as stratified by the value of the label. The random splits result in approximately (though not exactly) balanced problems. The same random samples are used for all quantum feature maps and hyperparameter values.

Kernel methods The quantum kernels are computed by using Qiskit²⁷. For each datapoint $\mathbf{x}_i$, we use the Qiskit Aer simulator²⁷ to compute the vector of amplitudes describing the quantum state of its representation $|\mathbf{x}_i\rangle$. Then, the values of the kernel $k(\mathbf{x}_i, \mathbf{x}_j) = |\langle \mathbf{x}_i | \mathbf{x}_j \rangle|^2$ are obtained by taking inner products of the corresponding vectors of amplitudes. The resulting kernel matrix is used in the scikit-learn²⁶ implementation of the SVC. Following⁷, we perform a grid search over the following values of the penalty term in Eq. 1:

$$C \in \{0.006, 0.015, 0.03, 0.0625, 0.125, 0.25, 0.5, 1.0, 2.0, 4.0, 8.0, 16.0, 32.0, 64.0, 128.0, 256, 512, 1024\}.$$

For the results presented in Fig. 1b, 5c, we choose the value C that maximizes the training score. In all other plots, we perform 5-fold cross-validation to choose C . In all plots, we report balanced accuracy, defined as the mean recall of the two classes, as the prediction accuracy score.

For the classical RBF kernel $k(\mathbf{x}_i, \mathbf{x}_j) = \exp(-\gamma \|\mathbf{x}_i - \mathbf{x}_j\|^2)$, we begin by performing 5-fold cross-validation over the values of C defined above and the following values of γ :

$$\gamma \in \{0.25, 0.5, 1.0, 2.0, 3.0, 4.0, 5.0, 20.0, 50.0, 100.0, 200.0, 500.0, 1000.0, 5000.0, 10000.0\} / (N \text{Var}[\mathbf{x}]),$$

where $\text{Var}[\mathbf{x}]$ is the variance of the training data and $N = 800$ is the number of training points. Then, we fix the value of C to the optimized value obtained from the cross-validation, and we vary parameter γ to obtain Fig. 5a.

Data availability We release all code used to generate the data and the plots, as well as the raw data, online: <https://github.com/rsln-s/Importance-of-Kernel-Bandwidth-in-Quantum-Machine-Learning/>.

References

1. Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Rev.* **41**, 303–332, DOI: [10.1137/s0036144598347011](https://doi.org/10.1137/s0036144598347011) (1999).
2. Schuld, M. Supervised quantum machine learning models are kernel methods (2021). [arXiv:2101.11020](https://arxiv.org/abs/2101.11020).
3. Belkin, M., Ma, S. & Mandal, S. To understand deep learning we need to understand kernel learning. In Dy, J. & Krause, A. (eds.) *Proceedings of the 35th International Conference on Machine Learning*, vol. 80 of *Proceedings of Machine Learning Research*, 541–549 (PMLR, 2018).
4. Schuld, M. & Killoran, N. Quantum machine learning in feature hilbert spaces. *Phys. Rev. Lett.* **122**, DOI: [10.1103/physrevlett.122.040504](https://doi.org/10.1103/physrevlett.122.040504) (2019).
5. Havlíček, V. *et al.* Supervised learning with quantum-enhanced feature spaces. *Nature* **567**, 209–212, DOI: [10.1038/s41586-019-0980-2](https://doi.org/10.1038/s41586-019-0980-2) (2019).
6. Liu, Y., Arunachalam, S. & Temme, K. A rigorous and robust quantum speed-up in supervised machine learning. *Nat. Phys.* **17**, 1013–1017, DOI: [10.1038/s41567-021-01287-z](https://doi.org/10.1038/s41567-021-01287-z) (2021).
7. Huang, H.-Y. *et al.* Power of data in quantum machine learning. *Nat. Commun.* **12**, DOI: [10.1038/s41467-021-22539-9](https://doi.org/10.1038/s41467-021-22539-9) (2021).
8. Arute, F. *et al.* Quantum supremacy using a programmable superconducting processor. *Nature* **574**, 505–510, DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5) (2019).
9. Bengio, Y., Delalleau, O. & Roux, N. L. The curse of highly variable functions for local kernel machines. In *Proceedings of the 18th International Conference on Neural Information Processing Systems*, NeurIPS’05, 107–114 (MIT Press, Cambridge, MA, USA, 2005).
10. Kübler, J. M., Buchholz, S. & Schölkopf, B. The inductive bias of quantum kernels. [arXiv:2106.03747](https://arxiv.org/abs/2106.03747) (2021).
11. Schölkopf, B., Smola, A. J., Bach, F. *et al.* *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond* (MIT press, 2002).
12. Hastie, T., Tibshirani, R. & Friedman, J. *The Elements of Statistical Learning* (Springer, 2009), second edn.
13. Vapnik, V. N. *The Nature of Statistical Learning* (Springer, 1996).
14. [Online] Retrieved 11/03/2021. https://www.tensorflow.org/quantum/tutorials/quantum_data.
15. Xiao, H., Rasul, K. & Vollgraf, R. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. [arXiv:1708.07747](https://arxiv.org/abs/1708.07747) (2017).
16. Clanuwat, T. *et al.* Deep learning for classical japanese literature. [arXiv:1812.01718](https://arxiv.org/abs/1812.01718) DOI: [10.20676/00000341](https://doi.org/10.20676/00000341) (2018).
17. The PLAsTiCC team *et al.* The photometric LSST astronomical time-series classification challenge (PLAsTiCC): Data set. [arXiv:1810.00001](https://arxiv.org/abs/1810.00001) (2018).
18. Peters, E. *et al.* Machine learning of high dimensional data on a noisy quantum processor. [arXiv:2101.09581](https://arxiv.org/abs/2101.09581) (2021).
19. Shawe-Taylor, J., Cristianini, N. *et al.* *Kernel Methods for Pattern Analysis* (Cambridge University Press, 2004).
20. Kung, S. Y. *Kernel Methods and Machine Learning* (Cambridge University Press, 2014).
21. Otten, M., Goumiri, I. R., Priest, B. W., Chapline, G. F. & Schneider, M. D. Quantum machine learning using gaussian processes with performant quantum kernels (2020). [arXiv:2004.11280](https://arxiv.org/abs/2004.11280).
22. [Online] https://github.com/rsln-s/Importance-of-Kernel-Bandwidth-in-Quantum-Machine-Learning/blob/c402b6ab2d761842e6ca198b9c577120e2af2e90/code/variance_depends_on_state.py.
23. Wiebe, N. & Kliuchnikov, V. Floating point representations in quantum circuit synthesis. *New J. Phys.* **15**, 093041, DOI: [10.1088/1367-2630/15/9/093041](https://doi.org/10.1088/1367-2630/15/9/093041) (2013).
24. [Online] Retrieved 11/03/2021. <https://ionq.com/best-practices>.

25. Lloyd, S., Schuld, M., Ijaz, A., Izaac, J. & Killoran, N. Quantum embeddings for machine learning. *arXiv:2001.03622* (2020).
26. Pedregosa, F. *et al.* Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* **12**, 2825–2830 (2011).
27. Abraham, H. *et al.* Qiskit: An open-source framework for quantum computing, DOI: [10.5281/zenodo.2573505](https://doi.org/10.5281/zenodo.2573505) (2021).

Acknowledgments

We thank Ilya Safro and Sami Khairy for helpful discussions about classical kernel methods and feedback on the manuscript. We thank Evan Peters for sharing the preprocessed dataset used in [18](#). This work was supported in part by the U.S. Department of Energy (DOE), Office of Science, Office of Advanced Scientific Computing Research AIDE-QC and FAR-QC projects and by the Argonne LDRD program under contract number DE-AC02-06CH11357. We gratefully acknowledge the computing resources provided on Bebop, a high-performance computing cluster operated by the Laboratory Computing Resource Center at Argonne National Laboratory.

Author contributions statement

R.S. conceived the experiments, developed the open-source code, and conducted the experiments. R.S. and S.M.W. contributed to technical discussions and writing of the manuscript.

Additional information

Competing interests

The authors declare no competing interests.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [supplementary.pdf](#)