

# QUACK: Quantum Aligned Centroid Kernel

Kilian Tscharke, Sebastian Issel, Pascal Debus

Quantum Security Technologies

Fraunhofer Institute for Applied and Integrated Security

Garching near Munich, Germany

$\langle\text{firstname}\rangle.\langle\text{lastname}\rangle@aisec.fraunhofer.de$

**Abstract**—Quantum computing (QC) seems to show potential for application in machine learning (ML). In particular quantum kernel methods (QKM) exhibit promising properties for use in supervised ML tasks. However, a major disadvantage of kernel methods is their unfavorable quadratic scaling with the number of training samples. Together with the limits imposed by currently available quantum hardware (NISQ devices) with their low qubit coherence times, small number of qubits, and high error rates, the use of QC in ML at an industrially relevant scale is currently impossible. As a small step in improving the potential applications of QKMs, we introduce QUACK, a quantum kernel algorithm whose time complexity scales linear with the number of samples during training, and independent of the number of training samples in the inference stage. In the training process, only the kernel entries for the samples and the centers of the classes are calculated, i.e. the maximum shape of the kernel for  $n$  samples and  $c$  classes is  $(n, c)$ . During training, the parameters of the quantum kernel and the positions of the centroids are optimized iteratively. In the inference stage, for every new sample the circuit is only evaluated for every centroid, i.e.  $c$  times. We show that the QUACK algorithm nevertheless provides satisfactory results and can perform at a similar level as classical kernel methods with quadratic scaling during training. In addition, our (simulated) algorithm is able to handle high-dimensional datasets such as MNIST with 784 features without any dimensionality reduction.

**Index Terms**—quantum computing, machine learning, kernel methods, linear complexity

## I. INTRODUCTION

Supervised Learning is an important branch of Machine Learning (ML) where a model is trained on labeled data to predict the labels of new, unseen data. It encompasses two main types of tasks: classification, which predicts discrete labels or classes, and regression, which forecasts continuous values. Quantum Machine Learning (QML) is an emerging field in the intersection of Quantum Computing (QC) and ML with the goal of utilizing the potential advantages of QC - like superposition, entanglement, and the exponential size of the Hilbert space - for Machine Learning. In particular, Quantum Kernel Methods (QKM) have recently gained attention because of their ability to replace many supervised quantum models and their guarantee to find equally good or better quantum models than variational circuits [1]. In addition, theoretical results are showing that QKMs can handle classification problems that cannot be solved using classical ML techniques, such as classifying numbers based on the discrete logarithm [2]. However, a significant drawback of using (quantum) kernels is the quadratic time complexity of the kernel calculation, i.e.  $\mathcal{O}(n_{\text{train}}^2)$ , since a kernel value must be estimated for each

pair of samples. For the inference stage, the time complexity without using advanced techniques such as Support Vectors is  $\mathcal{O}(n_{\text{train}}n_{\text{predict}})$ . Estimating a quantum kernel for the - by classical ML standards very small - URL dataset [3] with around 36,000 samples, requires approximately  $10^9$  kernel value calculations. With the commonly-used number of 1,000 shots, this involves  $10^{12}$  circuit executions. A state-of-the-art IBM device with Eagle r3 processor (e.g. *ibm\_sherbrooke* [4]) achieves 5,000 CLOPS (circuit executions per second) and hence the execution time for calculating the kernel would be  $10^8$  seconds, or over 6 years.

Quantum Kernel Alignment (QKA) is a fascinating tool for QKMs that uses kernels with variational parameters which can be trained to align the kernel to the ideal kernel for a given dataset [5]. This could enable the use of a general quantum kernel architecture that can be trained for different datasets.

The remainder of this paper is structured as follows: The next subsection I-A gives an overview of the related work in the fields of QKMs and QKA. The final part of the introduction contains our contributions (subsection I-B). In the following Background (section II), the fundamentals of supervised learning, quantum kernels, and QKA are introduced. Next, the Methods (section III) contain the implementation of the model and the Experiments (section IV) describe the numerical experiments carried out. In the Results and Discussion (section V) we show the results of our experiments and analyze them. Finally, Conclusion and Outlook (section VI) highlights the key results of this work and gives future research directions.

### A. Related Work

In 2021, Hubregetsen et al. [5] described the algorithm of QKA and defined the kernel-alignment measure. Moreover, they theoretically assessed the influence of noise on the algorithm and carried out numerical experiments on toy datasets, both on simulations and on real hardware.

Gentinetta et al. [6] developed a Quantum Support Vector Machine (QSVM) for which the quantum kernel is trained with QKA using the Pegasos algorithm in 2023. Unlike the default Support Vector Machine (SVM) implementation, their algorithm solves the primal formulation of the SVM which results in a min-min optimization and hence the SVM weights and the kernel parameters can be optimized simultaneously, increasing the efficiency of the algorithm.

In the same year, Kölle et al. [7] introduced a one-class QSVM, for which they reduced the training and inference times compared to a default QSVM by up to 95% and 25%, respectively, by applying randomized measurements and the variable subsampling ensemble method while achieving a superior average precision compared to a SVM with Radial Basis Function (RBF) kernel.

Finally, in 2024 Bowles et al. [8] benchmarked 12 popular QML models on six binary classification tasks. They concluded that out-of-the-box classical ML models tend to outperform the QML models and that entanglement does not necessarily improve the models' performance. Moreover, they noted that QML models both in simulations and on hardware can usually only handle input with size of the order of tens of features, and therefore classical pre-processing techniques such as principal component analysis are required to deal with higher dimensional data like the famous MNIST dataset. This classical pre-processing, however, influences the performance of the model and hence dilutes the results of a benchmark.

### B. Contributions

Our work aims to answer this question: Can the time complexity of QKMs be improved while still achieving satisfactory results?

We found a positive answer and report these contributions of our work:

- We develop Quantum Aligned Centroid Kernel (QUACK), a classifier based on quantum kernel alignment that improves the time complexity compared to basic kernel methods from  $\mathcal{O}(n_{\text{train}}^2)$  to  $\mathcal{O}(n_{\text{train}})$  during training and from  $\mathcal{O}(n_{\text{train}}n_{\text{test}})$  to  $\mathcal{O}(n_{\text{test}})$  during testing.
- We benchmark our classifier by evaluating it on eight different datasets with up to 784 features and different class ratios ranging from balanced to highly unbalanced.
- Finally, we observe that QUACK performs on a similar level as a classical SVM with RBF kernel

## II. BACKGROUND

### A. Supervised Learning

Let  $\mathcal{X} \subset \mathbb{R}^d$  be the data space and  $\mathbf{x} = (x_1, \dots, x_d) \in \mathcal{X}$  the feature vector of a single  $d$ -dimensional sample. Let  $\mathcal{Y}$  denote the target variable space and  $y \in \mathcal{Y}$  the target variable or label of a single sample. For the case of binary classification, we restrict the target variable to the set  $\{1, -1\}$ . Let further  $\Theta$  denote the space of model parameters. The general task of supervised machine learning is to train a parameterized model  $f_\theta : \mathcal{X} \times \Theta \rightarrow \mathcal{Y}$  such that it approximates a mapping between input  $\mathbf{x}$  and output  $\hat{y}$  based on the learned parameters  $\theta$ , as described in (1). During training, the parameters  $\theta$  are optimized such that the loss  $\mathcal{L}$  quantifying the difference between the predicted output  $\hat{y}$  and the target  $y$  is minimized, as in (2).

$$\hat{y} = f_\theta(\mathbf{x}) \quad (1)$$

$$\min_{\theta} \mathcal{L}(y, \hat{y}) \quad (2)$$

### B. Quantum Kernels

Quantum kernels emerge as an important tool for encoding classical data into quantum systems and subsequently classifying the data. It is hoped that the unique properties of quantum computing, such as entanglement and superposition, which are utilized in quantum kernels, will enable them to be more powerful than classical kernels. This hypothesis is supported by theoretical results showing that a constructed classification problem based on the discrete logarithm can be efficiently solved by QKMs, but not by classical ML methods [2]. In general, the encoding in QKMs is achieved through unitary operations  $U(\mathbf{x}_i)$  that depend on individual data points  $\mathbf{x}_i$ , often implemented via Pauli rotations. The state of the system after the encoding is

$$|\psi(\mathbf{x}_i)\rangle = |\psi_i\rangle = U(\mathbf{x}_i)|0\rangle. \quad (3)$$

Kernels are known from classical machine learning, where they are real- or complex-valued positive definite functions of two data points, i.e.  $\kappa : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{K}$ , where  $\mathbb{K} \in \{\mathbb{R}, \mathbb{C}\}$ . This definition can be extended to the quantum case, where a kernel  $k$  between two pure data-encoding quantum states  $\psi_i$  and  $\psi_j$  is calculated from the fidelity between these states

$$k(\mathbf{x}_i, \mathbf{x}_j) = F(\psi_i, \psi_j) = |\langle \psi_i | \psi_j \rangle|^2 \quad (4)$$

$$= |\langle 0^{\otimes n} | U^\dagger(\mathbf{x}_i) U(\mathbf{x}_j) | 0^{\otimes n} \rangle|^2 \quad (5)$$

with data encoding unitaries  $U(\mathbf{x}_j)$  and  $U^\dagger(\mathbf{x}_i)$ . This quantum kernel serves as a similarity measure between the states of two encoded samples: If both samples are identical, i.e.  $\mathbf{x}_i = \mathbf{x}_j$ , so  $\psi_i = \psi_j$  as well, the kernel equation (4) simplifies to

$$k(\mathbf{x}_i, \mathbf{x}_j) = k(\mathbf{x}_i, \mathbf{x}_i) = F(\psi_i, \psi_i) = |\langle \psi_i | \psi_i \rangle|^2 = 1. \quad (6)$$

On the other hand, if the encoded states  $\psi_i$  and  $\psi_j$  are orthogonal, the kernel will evaluate to

$$k(\mathbf{x}_i, \mathbf{x}_j) = F(\psi_i, \psi_j) = |\langle \psi_i | \psi_j \rangle|^2 = 0. \quad (7)$$

A quantum kernel can be implemented as an  $n$ -qubit circuit that consists of a trainable unitary  $U(\mathbf{x}_j)$ , encoding a single sample, followed by the inverse  $U^\dagger(\mathbf{x}_i)$  of another sample, and a measurement of all qubits, as shown in Fig. 1. The kernel value of the two samples is then obtained as the probability of measuring the all-zero state as given in equation (5). If a state vector simulator is used, the kernel value of two samples  $\mathbf{x}_i$  and  $\mathbf{x}_j$  is the fidelity of the states after application of the unitary  $U(\mathbf{x}_i)$ , respectively  $U(\mathbf{x}_j)$ , as given in (4).

### C. Trainable Quantum Kernels

A quantum kernel can contain not only parameters that encode the data into the circuit, but also adjustable parameters that affect the performance of the kernel for a particular dataset. This can for example be achieved by alternating layers of rotational gates, whose parameters consist of either one or more features of the datum  $\mathbf{x}$  or some other variational parameter  $\mathbf{w}$ . These parameters  $\mathbf{w}$  can be optimized through Quantum Kernel Alignment (QKA), as explained in subsection II-D. For variational circuits, *trainable encodings* seem to be

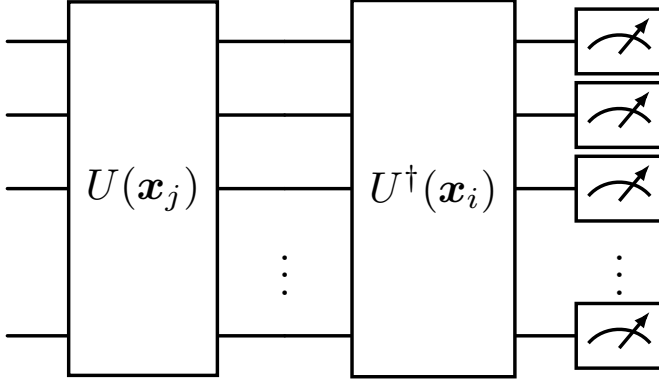


Fig. 1. Architecture of the circuit if executed on hardware. The kernel entry  $K_{ij}$  for samples  $i$  and  $j$  is the probability of measuring the all-zero bit string.

promising, since they were found to improve the robustness and generalization of the model [9]. A trainable encoding embeds the datum  $\mathbf{x}$  and the variational parameter  $\mathbf{w}$  and bias  $\mathbf{b}$  as one parameter vector  $\boldsymbol{\theta}$ , where each entry of  $\boldsymbol{\theta}$  is a single parameter used in a rotational gate. The parameter vector  $\boldsymbol{\theta}$  is calculated as

$$\boldsymbol{\theta} = \mathbf{w} \circ \mathbf{x} + \mathbf{b}, \quad (8)$$

analog to the neurons of neural networks, where  $\circ$  is the element-wise product (Hadamard product).

#### D. Quantum Kernel Alignment

QKA is a powerful tool that can be used to align a trainable kernel to the ideal kernel for a given dataset. It was originally developed for classical kernels [10] but can be used for quantum kernels as well. The implementation of QKA in this work is based on [5]. Kernel Alignment is used to optimize the kernel parameters for a specific task, improving the performance of kernel-based algorithms. The ideal kernel  $k^*$  is defined such that it always outputs the correct similarity between two data points:

$$k^*(\mathbf{x}_i, \mathbf{x}_j) = \begin{cases} 1 & \text{if } \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ in same class} \\ -1 & \text{if } \mathbf{x}_i \text{ and } \mathbf{x}_j \text{ in different classes} \end{cases} \quad (9)$$

In general, this ideal kernel is not known, but for the training set the ideal kernel matrix  $K^*$  can be constructed from the labels, i.e.  $K_{ij}^* = y_i y_j$ , or in vectorized form

$$K^* = \mathbf{y} \mathbf{y}^T. \quad (10)$$

The kernel-target alignment is a measure of the similarity between two kernels. To calculate it, we need the *Frobenius inner product* between two matrices as defined in (11).

$$\langle A, B \rangle_F = \sum_{ij} A_{ij} B_{ij} = \text{Tr} \{A^T B\} \quad (11)$$

With this, the kernel-target alignment TA between the current kernel  $K$  and the ideal kernel  $K^*$  can be calculated as in (12).

$$\begin{aligned} \text{TA}(K) &= \frac{\langle K, K^* \rangle_F}{\sqrt{\langle K, K \rangle_F \langle K^*, K^* \rangle_F}} \\ &= \frac{\sum_{ij} y_i y_j k(\mathbf{x}_i, \mathbf{x}_j)}{\sqrt{\left(\sum_{ij} k(\mathbf{x}_i, \mathbf{x}_j)^2\right) \left(\sum_{ij} y_i^2 y_j^2\right)}} \end{aligned} \quad (12)$$

The numerator of (12),  $\sum_{ij} y_i y_j k(\mathbf{x}_i, \mathbf{x}_j)$ , is the *kernel polarity*. If two samples are in the same class,  $y_i y_j = 1$ , the kernel value  $k(\mathbf{x}_i, \mathbf{x}_j)$  will increase the sum and hence the kernel-target alignment. For samples in different classes,  $y_i y_j = -1$ , the kernel polarity decreases by  $k(\mathbf{x}_i, \mathbf{x}_j)$  and subsequently the kernel-target alignment decreases, too. The kernel-target alignment equals 1 if the matrices are perfectly aligned and -1 if they are perfectly misaligned, i.e. perfectly inversely correlated.

### III. METHODOLOGY

Driven by the need of a NISQ compatible quantum classification algorithm that can handle data on an industrially relevant scale, we developed QUACK, a linear complexity algorithm for supervised classification based on quantum kernel alignment. QUACK is motivated by the desire to find a quantum kernel algorithm that avoids the calculation of the pairwise distances between the samples. Instead, our algorithm optimizes the distance using centroids as a proxy for each class with labels  $l \in \{1, -1\}$ . The centroids are initialized as the means of the classes in the original input data space. However, during training of the embedding map, the initial centroids cease to represent the center of the classes and we need to update them. This results in a two step alternating training procedure, where we iteratively optimize the parameters of the embedding map, followed by the position of one of the centroids. Since we do not want to store the centroids as a vector in the  $2^n$ -dimensional embedding space, we optimize the preimage of the centroids in embedding space, i.e. their positions in data space. Additionally, we alternate the class of the centroid to be optimized in each iteration.

The working principle of the algorithm is illustrated in Fig. 2. The centroids are initialized as the mean of each class in data space and initially, the embedding map performs a random embedding of the data in Hilbert space, as shown in the first figure. The first step of the algorithm, Kernel Alignment Optimization (KAO) iteration 1 for class 1 optimizes the parameters of the embedding map such that the distances between the samples of class 1 and centroid 1 are minimized, and the distances between the samples of class -1 and centroid 1 are maximized. This results in a new embedding map which is shown in the second figure. Next, the Centroid Optimization (CO) iteration 1 class -1 optimizes the position in data space of the centroid of the other class (class -1) with the aim of minimizing the distances between the centroid and the samples of class -1 and maximizing the distances between centroid -1 and the samples of class 1. The result of the CO step is

shown in the third figure. After this, the first epoch of the QUACK algorithm is complete, and a new iteration of the two step process starts. For the second iteration of the KAO, the embedding map is optimized such that the distances between the samples of class -1 and centroid -1 are minimal, and the distances between the samples of class 1 and centroid -1 are maximal. The resulting new embedding space is shown in the fourth figure. Next, the second iteration of the CO is carried out, where the position of centroid 1 is optimized in data space, followed by the next epoch of the two step process and so on.

The parameterized circuit with trainable encoding used for data encoding is described in subsubsection III-A1 Circuit Design and Data Encoding. The two-step optimization process is explained in subsubsection III-A2 Kernel Alignment Optimization and subsubsection III-A3 Centroid Optimization. Finally, the Prediction Stage - where the kernel entries of a new sample with the centroids are estimated and the sample is given the label of the class for whose centroid the kernel entry is maximal - is explained in subsubsection III-A4. The final part of this section, subsection III-B sketches very briefly how our state vector simulator works.

#### A. QUACK

The QUACK training algorithm is sketched in algorithm 1 and can be summarized as follows: The algorithm estimates a quantum kernel for the train samples  $X$  and the current working centroid  $c_l \in \{c_{-1}, c_1\}$ . Each of the  $n_{\text{epochs}}$  training epoch consists of a two-step optimization iterating between  $n_{\text{KAO}}$  epochs of optimizing the model parameters  $w, b$  and  $n_{\text{CO}}$  epochs of optimizing the centroids  $c_{-1}, c_1$ . For predicting new data, the kernel values of the new data  $X_{\text{predict}}$  and both centroids will be calculated, and each sample is given the label of the centroid with higher kernel entry. In the following, the different parts of the algorithm will be described in more detail.

---

#### Algorithm 1 QUACK training

---

**Input:** initial guess for  $c_{-1}, c_1$   
1:  $l \leftarrow \text{random bit} \cdot 2 - 1$   
2: **Repeat**  $n_{\text{epochs}}$  **times:**  
3:   **Repeat**  $n_{\text{KAO}}$  **times:**  
4:      $L_{\text{KAO}} \leftarrow \mathcal{L}_{\text{KAO}}(X, y, c_l, w, b)$   
5:     optimize model parameters  $w, b$   
6:    $l \leftarrow -l$   
7:   **Repeat**  $n_{\text{CO}}$  **times:**  
8:      $L_{\text{CO}} \leftarrow \mathcal{L}_{\text{CO}}(X, y, c_l, w, b)$   
9:     optimize  $c_l$

---



1) *Circuit Design and Data Encoding:* We use a trainable encoding map that was found to yield robustness and generalization improvements over fixed encodings [9]. In this context, trainable encoding refers to encodings, where the parameters of the gates depend on both, trainable weights and features of the data. How exactly the gate parameters are composed will be defined later.

The data encoding unitary  $U(x_j)$  consists of  $m' = m + 1$  layers of a unitary  $U_m(\theta_m)$ , as in Fig. 3. For clarification, if we

have e.g.  $m' = 5$  layers, the first layer is  $U_0(\theta_0)$  and the last layer is  $U_4(\theta_4)$ . Each of the unitaries  $U_m(\theta_m)$  is built using a layer of rotation gates and a ring of CNOT-gates, as sketched in Fig. 4. The rotation gate is the general parameterized rotation gate [11] with matrix representation shown in (13).

$$R(\theta_{m,1}, \theta_{m,2}, \theta_{m,3}) = R(\phi, \theta, \omega) = RZ(\omega)RY(\theta)RZ(\phi) = \begin{bmatrix} e^{-i(\phi+\omega)/2} \cos(\theta/2) & -e^{i(\phi-\omega)/2} \sin(\theta/2) \\ e^{-i(\phi-\omega)/2} \sin(\theta/2) & e^{i(\phi+\omega)/2} \cos(\theta/2) \end{bmatrix} \quad (13)$$

Each parameter  $\theta_{m,i}$  is calculated from the  $k$ -th feature of the sample  $x_j$  and weight  $w_{m,i}$  as given in (14), where  $k$  is a repeating counter from 1 to the number of input dimensions  $d$ , i.e.  $k = (3nm + i) \bmod d$  for the  $n$ -qubit system, the  $m$ -th layer and the  $i$ -th parameter in the layer.

$$\theta_{m,i} = w_{m,i} \cdot x_{j,k} + b_{m,i} \quad (14)$$

2) *Kernel Alignment Optimization:* During the Kernel Alignment Optimization, the parameters  $w$  and  $b$  of the embedding map are optimized. This is achieved by estimating the kernel between the train samples and one centroid and comparing it to the ideal kernel to obtain the kernel alignment. Since we use only one centroid to calculate the kernels, their matrices are  $n_{\text{samples}}$ -dimensional vectors. The ideal kernel is simply the label vector  $y$  if the centroid class is 1 and  $-y$  if the centroid class is -1. The current kernel entries are the fidelities between the current centroid  $c_l$  and the samples, where  $l \in \{-1, 1\}$  defines the class label of the current centroid:

$$k(c_l, x_i) = |\langle \psi_l | \psi_i \rangle|^2 \quad (15)$$

To get the kernel alignment, equation (12) is adapted for vectors as kernels and we obtain:

$$\text{TA} = \frac{l \cdot \sum_i y_i k(c_l, x_i)}{\sqrt{\left(\sum_i k(c_l, x_i)^2\right) \sum_i y_i^2}} \quad (16)$$

The loss function  $f_{c_l}$  is derived from the kernel-target alignment, with an additional regularization term with regularization parameter  $\lambda_{\text{KAO}}$ . Note that in this loss function, the centroid  $c_l$  is fixed.

$$f_{c_l}(w, b) = 1 - \text{TA} + \lambda_{\text{KAO}} \|w\|_2^2 \quad (17)$$

This loss function is then used to optimize the kernel parameters  $w$  and  $b$  either through backpropagation if the circuits are executed on a simulator or the parameter shift rule if real hardware is used, by solving this minimization problem:

$$\min_{w, b} f_{c_l}(w, b) \quad (18)$$

3) *Centroid Optimization:* The Centroid Optimization optimizes the position of the current centroid  $c_l$  in data space. For this, the kernel alignment is calculated the same way it is in the KAO optimization and then converted to a loss function  $g_{w,b}$ , in which the parameters  $w$  and  $b$  of the embedding map are fixed:

$$g_{w,b}(c_l) = 1 - \text{TA} + \lambda_{\text{CO}} R \quad (19)$$

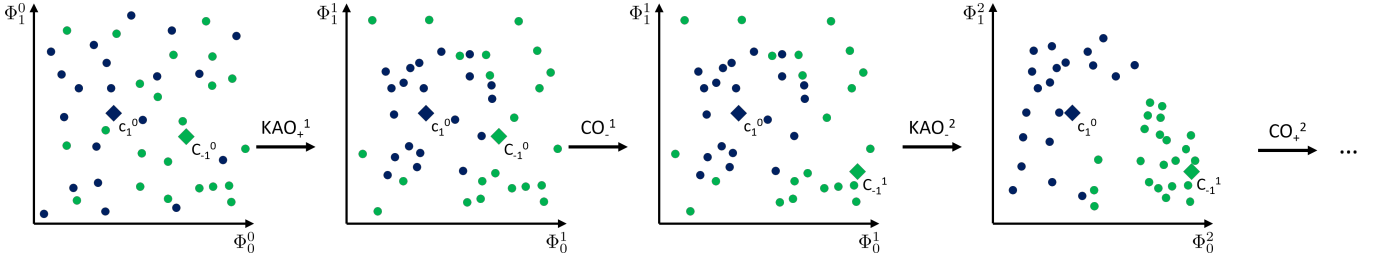


Fig. 2. Alternating optimization procedure for QUACK. The blue (green) dots show samples of class 1 (-1) in the embedding space  $\Phi$ . The superscript defines the current epoch and the subscript the dimension. The diamonds represent the centroids of the classes, where the superscript is the epoch in which the centroid has been optimized the last time and the subscript is the centroid class. For the KAO and CO steps, the superscript gives the epoch and the subscript the class for which the optimization is carried out with + representing class 1 and - representing class -1.

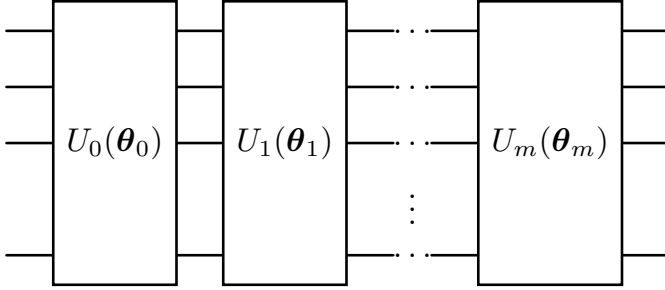


Fig. 3. Architecture of the encoding unitary  $U$ .

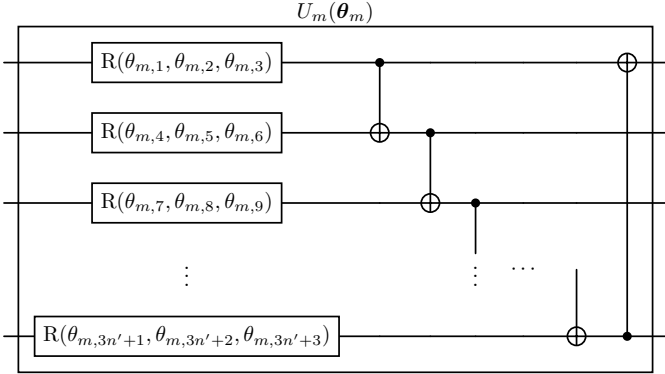


Fig. 4. Architecture of the unitary of a single parameterized layer  $U_m(\theta_m)$  with  $n' = n - 1$ .

Since the features are normalized in the range  $(0,1)$ , the regularization  $R$  introduces a penalty if the centroid position in data space is outside the normalization range:

$$R = \sum_d (\max(c_l^d - 1, 0) - \min(c_l^d, 0)), \quad (20)$$

where  $c_l^d$  is the  $d$ -th entry of the vector  $c_l$ . This loss is then used to optimize the position of the working centroid  $c_l$  in data space, by solving the following minimization problem:

$$\min_{c_l} g_{w,b}(c_l) \quad (21)$$

4) *Prediction Stage*: After the training is complete, the labels of new samples can be predicted. For this, the kernel

$K_{\text{predict}}$  of shape  $(n_{\text{samples}}, 2)$  between the new samples  $X_{\text{predict}}$  and both centroids  $c_{-1}$  and  $c_1$  is calculated. Each sample is given a label according to (22), where  $K_{i,1}$  represents the kernel value for sample  $i$  and centroid 1, and  $K_{i,-1}$  the value for sample  $i$  and centroid -1.

$$\hat{y}_i = \text{sign}(K_{i,1} - K_{i,-1}) \quad (22)$$

#### B. State Vector Simulator

To speed up the simulations, a state vector simulator is implemented using PyTorch [12]. The simulator builds up on the `nn.Module` class where the forward function returns the states of shape  $(n_{\text{samples}}, 2^{n_{\text{qubits}}})$  after applying all gates. The main advantage of the simulator is that the unitary of each gate is applied to all  $n_{\text{samples}}$  states in one operation, making the execution of the circuits for multiple samples dramatically faster when compared to other commonly used simulators.

### IV. EXPERIMENTS

Our linear time complexity QUACK algorithm is benchmarked on eight binary datasets from different areas and with various numbers of features and class ratios. The performance of our model is compared to three other models, containing both classical and quantum approaches. The code will be made available upon publication in a public code repository<sup>1</sup>.

#### A. Datasets

The model is benchmarked on eight datasets from different areas, including IT security and handwritten digits. An overview of the datasets is given in table I. The number of features varies between 14 for the Census dataset and 784 for the image datasets. The share of the smaller class in the total dataset varies between 0.09 (0.07/0.09) and 0.50 (0.49/0.49) for the train (validation/test) set, meaning there are both balanced and highly unbalanced datasets. For the datasets that have not been pre-split into test and train sets, the train (test) set is created by randomly selecting 70% (30%) of the samples from the dataset. Next, 1,000 samples are randomly selected from the training set for training, 400 from the test set for validation, and a further 400 from the test set for the final testing. The class labels are  $\{1, -1\}$  according to the criteria specified in table I.

<sup>1</sup><https://github.com/kilian106/QUACK>

TABLE I

OVERVIEW OF THE DATASETS USED IN THE BENCHMARK. THE RATIOS SHOW THE RATIO OF THE MINORITY CLASS TO THE NUMBER OF SAMPLES IN THE SET.

| Dataset | Ref. | Description         | Class 1    | Class -1   | Ratio Train | Ratio Val. | Ratio Test | Features |
|---------|------|---------------------|------------|------------|-------------|------------|------------|----------|
| Census  | [13] | Income              | $\leq 50K$ | $> 50K$    | 0.28        | 0.26       | 0.23       | 14       |
| CoverT  | [14] | Forest tree types   | 4          | $> 4$      | 0.09        | 0.07       | 0.09       | 15       |
| DoH     | [15] | Network traffic     | Benign     | Malicious  | 0.24        | 0.23       | 0.21       | 33       |
| EMNIST  | [16] | Handwritten letters | A-M        | N-Z        | 0.50        | 0.48       | 0.48       | 784      |
| FMNIST  | [17] | Clothing types      | 0-4        | 5-9        | 0.50        | 0.49       | 0.49       | 784      |
| KDD     | [18] | Network intrusion   | Normal     | Anomalous  | 0.50        | 0.44       | 0.47       | 42       |
| MNIST   | [19] | Handwritten digits  | 0-4        | 5-9        | 0.48        | 0.47       | 0.49       | 784      |
| URL     | [20] | URLs                | Benign     | Non-benign | 0.15        | 0.14       | 0.15       | 79       |

### B. Models for Benchmarking

For a comprehensive evaluation of our model, it is benchmarked against these three other approaches: 1) A *vanilla SVM with RBF kernel* and default parameters, implemented with scikit-learn [21]. This model has a quadratic time complexity during training and a linear during testing 2) An *RBF centroid classifier*. This classifier first determines the mean of each class in feature space and then calculates a RBF kernel of shape  $(n_{\text{samples}}, 2)$  that contains the kernel values between the samples and the means of both classes. Each sample is given the label of the class for which the kernel value between the sample and the respective centroid is larger. This classifier has a linear time complexity and does not require any training. 3) A *Quantum Support Vector Machine* that uses the trained kernel parameters from our approach. However, the QSVM determines the kernels in their quadratic form  $(n_{\text{train}}, n_{\text{train}})$  for the training of the SVM parameters and of the shape  $(n_{\text{test}}, n_{\text{train}})$  for testing. The SVM hyperparameters are the default ones from scikit-learn.

### C. Implementation Details

All models are trained three times with different random seeds and the mean results with standard deviation are reported. The hyperparameters used for the QUACK algorithm are shown in tables II and III in Appendix A. The hyperparameter tuning is achieved by a randomized grid search for each dataset where the validation set is used to evaluate the model. The number of layers and qubits were selected such that each feature is encoded at least once and the model can still be simulated in a reasonable time.

### D. Verification of the Simulator

To make sure that the state vector simulator works as intended, a small model is trained on both our simulator and PennyLane’s [22] *default.qubit* simulator with identical hyperparameters. After training, the weights, biases, loss, and metrics of both models were identical, and therefore we can conclude that our simulator works as intended.

## V. RESULTS AND DISCUSSION

The newly introduced linear time complexity algorithm QUACK is benchmarked together with three other models on eight binary datasets from different areas with various numbers of features and class ratios. Each model is run three times on each dataset with different random seeds.

### A. Model Performance

Figure 5 shows the average of the test area under the ROC curve (AUC) for each model and dataset, and the values are listed in table IV in Appendix B. Our model performs equally or almost equally as the SVM RBF on five out of eight datasets. More precisely, for CoverT, DoH, FMNIST, KDD, and MNIST, the difference in AUC is 0.02 at most. The performance gap is highest on EMNIST with an AUC difference of 0.06, followed by Census and URL with 0.03. Fig. 7 shows the test AUCs of the best run of each model. The main difference to Fig. 5 is that the best QUACK run additionally achieves equal results as the SVM on the Census dataset. From this, we conclude that QUACK performs on a similar level as the SVM and may be a reasonable alternative to the SVM.

The QSVM that uses the trained weights from our model to compute the full kernel, achieves very similar AUCs compared to our model, with the highest difference being 0.02 in both directions. This suggests, that once the kernel training is completed and the kernel parameters are set, the SVM training and inference methods do not notably improve the model’s performance.

The RBF centroid classifier is the worst model on all datasets, which is intuitive since this model does not require any training at all. It is surprising, however, that this classifier comes relatively close to the performance of the other models for the DoH and KDD datasets. Together with the particularly good performance of the other models on these two datasets, we suspect that DoH and KDD are relatively easy datasets for binary classification.

Finally, the observed standard deviation for QUACK across datasets is consistently low, being 0.01 or below, except for Census and CoverT. These two datasets are notable outliers with a standard deviation of 0.04 and 0.02, respectively. From this, we conclude that the performance of QUACK is largely independent of the initialization of the trainable parameters. The QSVM shows a similar standard deviation as QUACK which was expected since both algorithms use the same optimized weights and biases. The SVM RBF and RBF Centroid exhibit no standard deviation, as they are deterministic algorithms.

### B. Number of Circuit Evaluations

We compare the number of circuit evaluations required during the training of our model with a standard kernel method. All circuit evaluations that result from an evaluation of the models

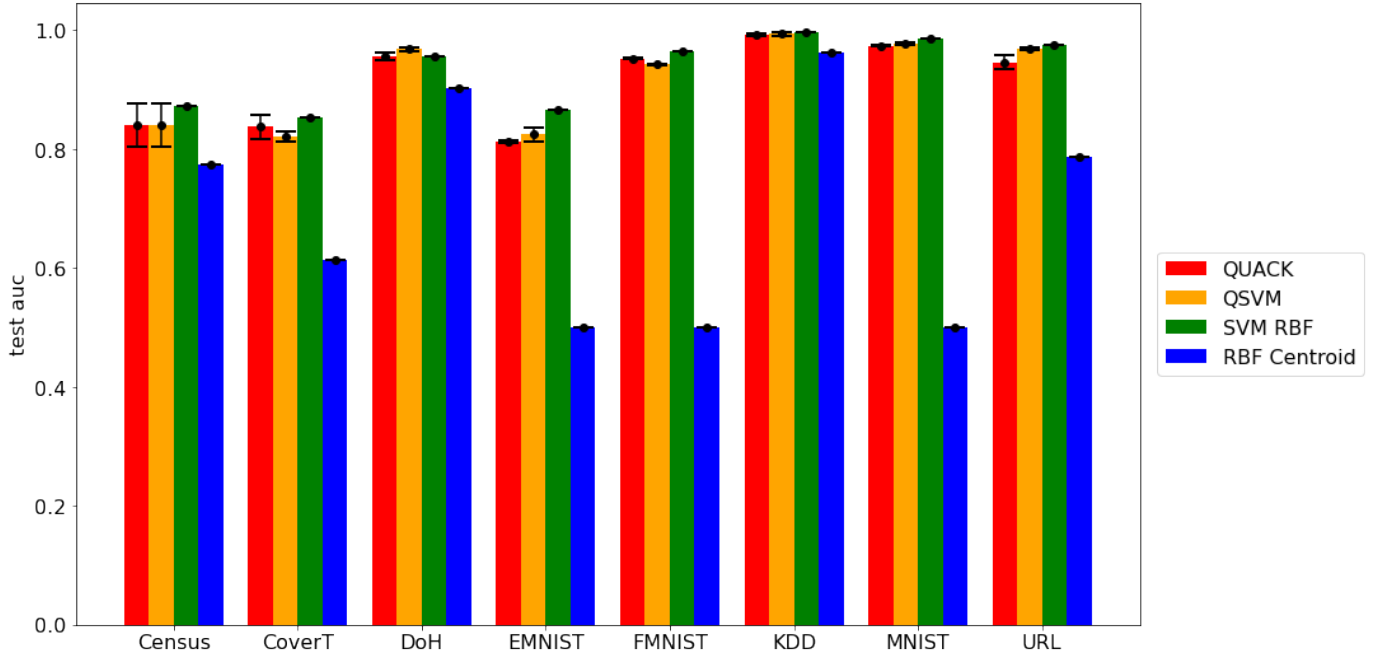


Fig. 5. Test AUCs of the different models.

during training are ignored. Fig. 6 shows the number of circuit evaluations over the number of train samples for the number of epochs used throughout the numerical experiments (see table III in Appendix A). The QUACK algorithm scales linearly with the number of training samples  $n_{\text{train}}$ , and the number of circuit evaluations is  $N_{\text{QUACK}} = n_{\text{epochs}} \cdot (n_{\text{KAO}} + n_{\text{CO}}) \cdot n_{\text{train}}$ , where  $n_{\text{epochs}}$  is the number of two-step iterations performed,  $n_{\text{KAO}}$  and  $n_{\text{CO}}$  are the numbers of the Kernel Alignment Optimization steps and Centroid Optimization steps, respectively. A standard kernel on the other hand, requires a quadratic number of circuit evaluations  $N_{\text{standard kernel}} = n_{\text{epochs}} \cdot n_{\text{train}}^2$ . As soon as the number of samples exceeds the sum of the number of epochs for kernel alignment and centroid optimization, i.e.  $n_{\text{train}} > n_{\text{KAO}} + n_{\text{CO}}$ , QUACK needs fewer circuit evaluations than the default kernel. With a further increase in the sample size, the number of circuit evaluations required for QUACK grows quadratically slower than for the standard kernel.

## VI. CONCLUSION AND OUTLOOK

We developed QUACK, a classifier based on quantum kernel alignment that improves the time complexity compared to basic kernel methods from  $\mathcal{O}(n_{\text{train}}^2)$  to  $\mathcal{O}(n_{\text{train}})$  during training and from  $\mathcal{O}(n_{\text{train}}n_{\text{test}})$  to  $\mathcal{O}(n_{\text{test}})$  during testing. QUACK's training time complexity is a polynomial improvement compared to the SVM. The algorithm was benchmarked by evaluating it on eight different datasets with up to 784 features and various class ratios ranging from balanced to highly unbalanced. The performance was compared to a vanilla SVM with an RBF kernel, an RBF centroid classifier, and a QSVM. We conclude that QUACK performs on a similar level as the classical SVM and that the training of the SVM parameters does not improve the

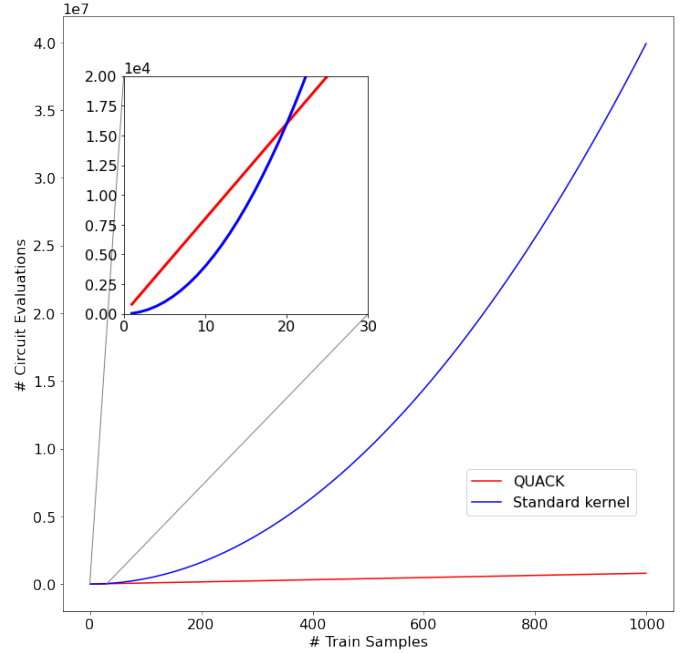


Fig. 6. Comparison of the number of evaluation steps between QUACK and a standard kernel. The inset shows a zoom-in of the plot for the number of train samples in the range from 0 to 30.

predictions of the model once the kernel parameters are trained. Finally, our algorithm works on data with up to 784 features without dimensionality reduction, which is often required for other state-of-the-art QML models.

Thanks to the linear scaling of QUACK, the algorithm

can be used as quick baseline for future classification tasks: If the performance of QUACK is satisfactory, using more costly classification algorithms does not offer an advantage. If, however, QUACK performs poorly, the application of more costly algorithms, e.g. the quadratic-scaling (Q)SVM, should be considered. Furthermore, there is an intuitive explanation for the classification results of the algorithm: If QUACK performs well, it has found an encoding circuit and centroids which separate the classes into different clusters around these centroids in Hilbert space.

On the other hand, QUACK is limited by the assumption that there exists an embedding in which each class forms a cluster around a different centroid. Finally, the performance of QUACK is heavily dependent on the choice of hyperparameters, and a extensive tuning of these hyperparameters is strongly recommended.

This work is only a first effort toward increasing the potential of QKMs and the next logical step is to benchmark QUACK on hardware. Further work can extend the algorithm to perform multi-class classification by using  $k$  centroids for  $k$  classes and running a one-versus-all classification in each iteration of the algorithm. In addition, the stability of the algorithm should be improved to make its performance less dependent on the choice of hyperparameters.

#### ACKNOWLEDGMENT

This research is part of the Munich Quantum Valley, which is supported by the Bavarian state government with funds from the Hightech Agenda Bayern Plus.

#### REFERENCES

- [1] M. Schuld, *Supervised quantum machine learning models are kernel methods*, 2021. arXiv: 2101.11020 [quant-ph].
- [2] Y. Liu, S. Arunachalam, and K. Temme, “A rigorous and robust quantum speed-up in supervised machine learning,” *Nature Physics*, vol. 17, pp. 1013–1017, 9 Sep. 2021, ISSN: 17452481. DOI: 10.1038/s41567-021-01287-z.
- [3] M. S. I. Mamun, M. A. Rathore, A. H. Lashkari, N. Stakhanova, and A. A. Ghorbani, “Detecting malicious urls using lexical analysis,” in *Network and System Security*, J. Chen, V. Piuri, C. Su, and M. Yung, Eds., Cham: Springer International Publishing, 2016, pp. 467–482, ISBN: 978-3-319-46298-1.
- [4] *Ibm\_sherbrooke*, [https://quantum.ibm.com/services/resources?system=ibm\\_sherbrooke](https://quantum.ibm.com/services/resources?system=ibm_sherbrooke), Accessed: 2024-03-28.
- [5] T. Hubregtsen *et al.*, “Training quantum embedding kernels on near-term quantum computers,” *Phys. Rev. A*, vol. 106, p. 042431, 4 2022. DOI: 10.1103/PhysRevA.106.042431. [Online]. Available: <https://link.aps.org/doi/10.1103/PhysRevA.106.042431>.
- [6] G. Gentinetta, D. Sutter, C. Zoufal, B. Fuller, and S. Woerner, “Quantum kernel alignment with stochastic gradient descent,” in *2023 IEEE International Conference on Quantum Computing and Engineering (QCE)*, IEEE, Sep. 2023. DOI: 10.1109/qce57702.2023.00036. [Online]. Available: <http://dx.doi.org/10.1109/QCE57702.2023.00036>.
- [7] M. Kölle *et al.*, *Towards efficient quantum anomaly detection: One-class svms using variable subsampling and randomized measurements*, 2023. arXiv: 2312.09174 [quant-ph].
- [8] J. Bowles, S. Ahmed, and M. Schuld, *Better than classical? the subtle art of benchmarking quantum machine learning models*, 2024. arXiv: 2403.07059 [quant-ph].
- [9] J. Berberich, D. Fink, D. Pranjic, C. Tutschku, and C. Holm, *Training robust and generalizable quantum models*, 2023. arXiv: 2311.11871 [quant-ph].
- [10] N. Cristianini, J. Shawe-Taylor, A. Elisseeff, and J. Kandola, “On kernel-target alignment,” in *Advances in Neural Information Processing Systems*, T. Dietterich, S. Becker, and Z. Ghahramani, Eds., vol. 14, MIT Press, 2001. [Online]. Available: [https://proceedings.neurips.cc/paper\\_files/paper/2001/file/1f71e393b3809197ed66df836fe833e5-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2001/file/1f71e393b3809197ed66df836fe833e5-Paper.pdf).
- [11] M. Schuld and F. Petruccione, “Quantum computing,” in *Machine Learning with Quantum Computers*. Cham: Springer International Publishing, 2021, pp. 79–146, ISBN: 978-3-030-83098-4. DOI: 10.1007/978-3-030-83098-4\_3. [Online]. Available: [https://doi.org/10.1007/978-3-030-83098-4\\_3](https://doi.org/10.1007/978-3-030-83098-4_3).
- [12] A. Paszke *et al.*, *Pytorch: An imperative style, high-performance deep learning library*, 2019. arXiv: 1912.01703 [cs.LG].
- [13] D. Dua and C. Graff, *UCI machine learning repository*, 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>.
- [14] J. A. Blackard and D. J. Dean, “Comparative accuracies of artificial neural networks and discriminant analysis in predicting forest cover types from cartographic variables,” *Computers and Electronics in Agriculture*, vol. 24, no. 3, pp. 131–151, 1999, ISSN: 0168-1699. DOI: [https://doi.org/10.1016/S0168-1699\(99\)00046-0](https://doi.org/10.1016/S0168-1699(99)00046-0). [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S0168169999000460>.
- [15] M. MontazeriShatoori, L. Davidson, G. Kaur, and A. Habibi Lashkari, “Detection of doh tunnels using time-series classification of encrypted traffic,” in *2020 IEEE Intl Conf on Dependable, Autonomic and Secure Computing, Intl Conf on Pervasive Intelligence and Computing, Intl Conf on Cloud and Big Data Computing, Intl Conf on Cyber Science and Technology Congress (DASC/PiCom/CBDCCom/CyberSciTech)*, 2020, pp. 63–70. DOI: 10.1109/DASC-PiCom-CBDCCom-CyberSciTech49142.2020.00026.
- [16] G. Cohen, S. Afshar, J. Tapson, and A. van Schaik, “Emnist: Extending mnist to handwritten letters,” in

2017 *International Joint Conference on Neural Networks (IJCNN)*, 2017, pp. 2921–2926. DOI: 10.1109/IJCNN.2017.7966217.

- [17] H. Xiao, K. Rasul, and R. Vollgraf, *Fashion-mnist: A novel image dataset for benchmarking machine learning algorithms*, 2017. arXiv: 1708.07747 [cs.LG].
- [18] M. Tavallaei, E. Bagheri, W. Lu, and A. A. Ghorbani, “A detailed analysis of the kdd cup 99 data set,” in *2009 IEEE Symposium on Computational Intelligence for Security and Defense Applications*, 2009, pp. 1–6. DOI: 10.1109/CISDA.2009.5356528.
- [19] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner, “Gradient-based learning applied to document recognition,” *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998. DOI: 10.1109/5.726791.
- [20] M. S. I. Mamun, M. A. Rathore, A. H. Lashkari, N. Stakhanova, and A. A. Ghorbani, “Detecting malicious urls using lexical analysis,” in *Network and System Security*, J. Chen, V. Piuri, C. Su, and M. Yung, Eds., Cham: Springer International Publishing, 2016, pp. 467–482, ISBN: 978-3-319-46298-1.
- [21] F. Pedregosa *et al.*, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [22] V. Bergholm *et al.*, *Pennylane: Automatic differentiation of hybrid quantum-classical computations*, 2022. arXiv: 1811.04968 [quant-ph].

## APPENDIX A HYPERPARAMETERS

TABLE II  
OVERVIEW OF THE OPTIMIZED HYPERPARAMETERS FOR EACH DATASET.

| Datasets | $lr_{kao}$ | $lr_{co}$ | $r_{decay}$ | $reg_{kao}$ | $reg_{co}$ |
|----------|------------|-----------|-------------|-------------|------------|
| Census   | 0.5        | 0.5       | 0.9         | 0.001       | 0.001      |
| CoverT   | 0.5        | 0.1       | 0.9         | 0.001       | 0.001      |
| DoH      | 0.5        | 0.5       | 0.9         | 0.001       | 0.001      |
| EMNIST   | 1.0        | 5.0       | 0.9         | 0.001       | 0.001      |
| FMNIST   | 5.0        | 0.5       | 0.8         | 0.0001      | 0.001      |
| KDD      | 0.5        | 1.0       | 0.9         | 0.001       | 0.001      |
| MNIST    | 5.0        | 1.0       | 0.9         | 0.001       | 0.001      |
| URL      | 0.5        | 0.5       | 0.9         | 0.001       | 0.001      |

TABLE III  
OVERVIEW OF THE HYPERPARAMETERS SHARED BETWEEN QUACK ON ALL DATASETS. THE NUMBER OF EPOCHS FOR THE TWO-STEP TRAINING, KERNEL ALIGNMENT OPTIMIZATION AND CENTROID OPTIMIZATION ARE GIVEN BY  $n$ ,  $n_{kao}$ , AND  $n_{co}$  RESPECTIVELY. INIT\_WEIGHTS\_SCALE GIVES THE MAXIMUM VALUE FOR THE WEIGHTS DURING RANDOM INITIALIZATION.

| layers | qubits | $n_{train}$ | $n_{val}$ | $n_{test}$ | $n$ | $n_{kao}$ | $n_{co}$ | init_weights_scale | seeds         |
|--------|--------|-------------|-----------|------------|-----|-----------|----------|--------------------|---------------|
| 53     | 5      | 1000        | 400       | 400        | 40  | 10        | 10       | 0.1                | 42, 123, 1234 |

## APPENDIX B DETAILED RESULTS

TABLE IV  
OVERVIEW OF THE AUCS OF THE MODELS.

| Dataset | train_auc       | val_auc         | test_auc        | qsvm_train_auc  | qsvm_val_auc    | qsvm_test_auc   |
|---------|-----------------|-----------------|-----------------|-----------------|-----------------|-----------------|
| Census  | $0.91 \pm 0.03$ | $0.85 \pm 0.02$ | $0.84 \pm 0.04$ | $0.91 \pm 0.03$ | $0.85 \pm 0.02$ | $0.84 \pm 0.04$ |
| CoverT  | $0.85 \pm 0.03$ | $0.75 \pm 0.04$ | $0.84 \pm 0.02$ | $0.88 \pm 0.02$ | $0.73 \pm 0.05$ | $0.82 \pm 0.01$ |
| DoH     | $0.98 \pm 0.00$ | $0.96 \pm 0.00$ | $0.96 \pm 0.01$ | $0.98 \pm 0.00$ | $0.98 \pm 0.00$ | $0.97 \pm 0.00$ |
| EMNIST  | $0.99 \pm 0.01$ | $0.84 \pm 0.01$ | $0.81 \pm 0.00$ | $0.99 \pm 0.00$ | $0.84 \pm 0.01$ | $0.82 \pm 0.01$ |
| FMNIST  | $0.98 \pm 0.00$ | $0.96 \pm 0.00$ | $0.95 \pm 0.00$ | $0.99 \pm 0.00$ | $0.95 \pm 0.01$ | $0.94 \pm 0.00$ |
| KDD     | $1.00 \pm 0.00$ | $0.99 \pm 0.00$ | $0.99 \pm 0.00$ | $1.00 \pm 0.00$ | $0.99 \pm 0.00$ | $0.99 \pm 0.00$ |
| MNIST   | $0.99 \pm 0.00$ | $0.95 \pm 0.01$ | $0.97 \pm 0.00$ | $1.00 \pm 0.00$ | $0.96 \pm 0.01$ | $0.98 \pm 0.00$ |
| URL     | $0.93 \pm 0.01$ | $0.89 \pm 0.02$ | $0.95 \pm 0.01$ | $0.95 \pm 0.01$ | $0.92 \pm 0.01$ | $0.97 \pm 0.00$ |

| Dataset | svm_rbf_train_auc | svm_rbf_val_auc | svm_rbf_test_auc | rbf_centroid_val_auc | rbf_centroid_test_auc |
|---------|-------------------|-----------------|------------------|----------------------|-----------------------|
| Census  | $0.88 \pm 0.00$   | $0.87 \pm 0.00$ | $0.87 \pm 0.00$  | $0.73 \pm 0.00$      | $0.77 \pm 0.00$       |
| CoverT  | $0.92 \pm 0.00$   | $0.79 \pm 0.00$ | $0.85 \pm 0.00$  | $0.63 \pm 0.00$      | $0.61 \pm 0.00$       |
| DoH     | $0.98 \pm 0.00$   | $0.96 \pm 0.00$ | $0.96 \pm 0.00$  | $0.90 \pm 0.00$      | $0.90 \pm 0.00$       |
| EMNIST  | $0.99 \pm 0.00$   | $0.89 \pm 0.00$ | $0.87 \pm 0.00$  | $0.50 \pm 0.00$      | $0.50 \pm 0.00$       |
| FMNIST  | $0.99 \pm 0.00$   | $0.97 \pm 0.00$ | $0.97 \pm 0.00$  | $0.50 \pm 0.00$      | $0.50 \pm 0.00$       |
| KDD     | $1.00 \pm 0.00$   | $1.00 \pm 0.00$ | $1.00 \pm 0.00$  | $0.95 \pm 0.00$      | $0.96 \pm 0.00$       |
| MNIST   | $1.00 \pm 0.00$   | $0.98 \pm 0.00$ | $0.98 \pm 0.00$  | $0.50 \pm 0.00$      | $0.50 \pm 0.00$       |
| URL     | $0.97 \pm 0.00$   | $0.95 \pm 0.00$ | $0.98 \pm 0.00$  | $0.71 \pm 0.00$      | $0.79 \pm 0.00$       |

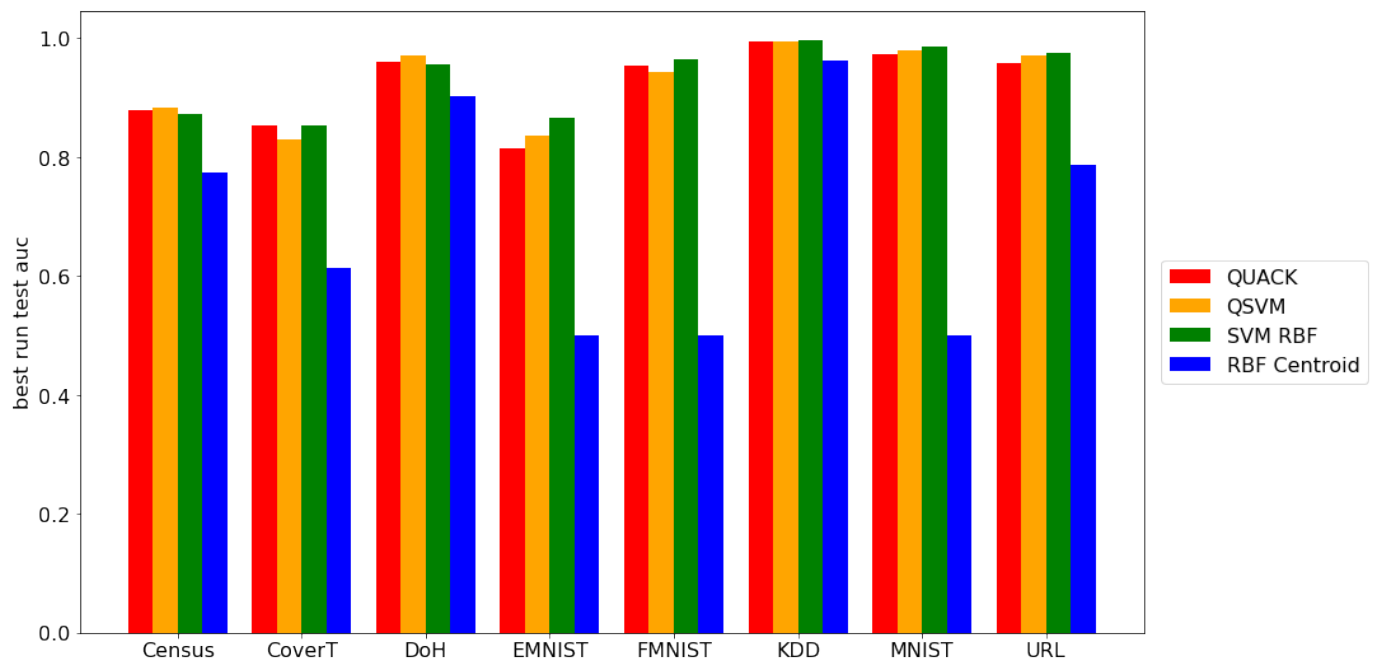


Fig. 7. Test AUCs of the best run of each model for each dataset.