



Quantum Advantage Seeker with Kernels (QuASK): a software framework to speed up the research in quantum machine learning

Francesco Di Marcantonio¹ · Massimiliano Incudini^{2,4} · Davide Tezza^{3,4} · Michele Grossi¹

Received: 30 June 2022 / Accepted: 8 March 2023
© The Author(s) 2023

Abstract

Exploiting the properties of quantum information to the benefit of machine learning models is perhaps the most active field of research in quantum computation. This interest has supported the development of a multitude of software frameworks (e.g. Qiskit, PennyLane, Braket) to implement, simulate, and execute quantum algorithms. Most of them allow us to define quantum circuits, run basic quantum algorithms, and access low-level primitives depending on the hardware such software is supposed to run. For most experiments, these frameworks have to be manually integrated within a larger machine learning software pipeline. The researcher is in charge of knowing different software packages, integrating them through the development of long code scripts, analyzing the results, and generating the plots. Long code often leads to erroneous applications, due to the average number of bugs growing proportional with respect to the program length. Moreover, other researchers will struggle to understand and reproduce the experiment, due to the need to be familiar with all the different software frameworks involved in the code script. We propose QuASK, an open-source quantum machine learning framework written in Python that aids the researcher in performing their experiments, with particular attention to quantum kernel techniques. QuASK can be used as a command-line tool to download datasets, pre-process them, quantum machine learning routines, analyze and visualize the results. QuASK implements most state-of-the-art algorithms to analyze the data through quantum kernels, with the possibility to use projected kernels, (gradient-descent) trainable quantum kernels, and structure-optimized quantum kernels. Our framework can also be used as a library and integrated into pre-existing software, maximizing code reuse.

Keywords Software for quantum machine learning · Software · Quantum kernels · Quantum machine learning · Quantum computing

1 Introduction

Breakthroughs in quantum technologies have allowed the construction of small-scale prototypes of quantum

computers (Madsen et al. 2022; Dumitrescu et al. 2022; Huang et al. 2022), namely NISQ devices (Preskill 2018). Even though many sources of noise may corrupt the execution on these devices (Pelofske et al. 2022), we are able to run a certain class of algorithms (Bharti et al. 2022) which compromises the strong theoretical speedup of fault-tolerant quantum algorithms (Montanaro 2016) to achieve shorter, less noisy computations. A large subset of the NISQ-ready algorithms is dedicated to the development of machine learning models.

One of the most interesting technique among them are the quantum classifiers (Schuld and Killoran 2019; Havlíček et al. 2019; Mengoni and Di Pierro 2019): the function $f(x) = \text{Tr}[\rho_x \rho_w]$, where ρ_x represents the encoding of a data point \mathbf{x} in a quantum state through the feature map $|0\rangle\langle 0| \mapsto U(\mathbf{x})|0\rangle\langle 0| = \rho_x$ and ρ_w represents the weight vector encoded through the mapping $|0\rangle\langle 0| \mapsto$

✉ Massimiliano Incudini
massimiliano.incudini@univr.it

✉ Michele Grossi
michele.grossi@cern.ch

¹ European Organization for Nuclear Research (CERN),
Geneva 1211, Switzerland

² Dipartimento di Informatica, Università di Verona, Verona
34137, Italy

³ Dipartimento di Matematica, Università di Trento, Povo
38123, Italy

⁴ Data Reply s.r.l., Turin 10126, Italy

$W|0\rangle\langle 0| = \rho_w$, can be interpreted as a linear¹ model (Schuld and Petruccione 2021). Such a function can be immediately used to solve supervised learning tasks. By choosing the weight mapping to be parametric $W(\theta)$, we can train the parameters to minimize some loss function using gradient-descent-based techniques: such an approach is named *quantum neural network* (Mitarai et al. 2018; Abbas et al. 2021). However, the training phase of these models could be affected by barren plateau (McClean et al. 2018; Holmes et al. 2021), i.e. the flat loss landscape, where the variance of the gradient vanishing exponentially fast with respect to the number of qubits. Highly entangled states (Marrero et al. 2021), noise (Wang et al. 2021), global measurement (Arrasmith et al. 2021), and expressibility of the feature map (Holmes et al. 2022) have been linked to the appearance of barren plateau. To avoid such a problem, *Quantum Kernel Estimation* (QKE) (Schuld and Killoran 2019; Havlíček et al. 2019; Mengoni and Di Pierro 2019) algorithm can be used in a hybrid form - we implement a quantum kernel function $\kappa(\mathbf{x}, \mathbf{x}') = \text{Tr}[\rho_{\mathbf{x}}\rho_{\mathbf{x}'}]$, quantifying the similarity between two encoded data points, with a classical machine learning algorithm. The training of the model is classical and is expected to end successfully² and efficiently due to the representer theorem (Schölkopf et al. 2001). Classical kernel methods are a cornerstone of machine learning, and have been applied to any sort of task including signal processing (Pérez-Cruz and Bousquet 2004; Rojo-Álvarez et al. 2018; Camps-Valls 2006), bioinformatics (Camps-Valls 2006; Ben-Hur et al. 2008), and image processing (Wang and Qi 2014; Yang 2001).

A clear benefit in using quantum kernel estimation to enhance machine learning applications has still to be found (Schuld and Killoran 2022). Quantum kernels have been shown to improve the performances of classical machine learning algorithms for some problems, such as the prediction of the output of quantum systems (Huang et al. 2021, 2022), and in learning from distributions based on the discrete logarithm (Liu et al. 2021). They have been applied to several real-world, industrial scale problems such as anomaly detection (Liu and Rebentrost 2018), fraud detection (Di Pierro and Incudini 2021; Grossi et al. 2022; Kyriienko and Magnusson 2022), the effectiveness of pharmaceutical treatments (Krunić et al. 2022), and supernova classifications (Peters et al. 2021). These approaches have been experimentally tested on superconducting (Peters et al. 2021; Wang

et al. 2021), optical (Bartkiewicz et al. 2020), and NMR (Kusumoto et al. 2021) quantum devices, and their effectiveness is usually assessed empirically.

Most of these experiments share, at least partially, a common structure: dimensionality reduction techniques, used to limit the number of quantum resources needed for the computation; the scaling of the input; the choice of the quantum kernel; the evaluation of the quantum kernel. The researcher is usually in charge of developing a software prototype, which requires the knowledge of many different software frameworks and platforms: the ones dealing with the machine learning tasks (Paszke et al. 2019; Chollet et al. 2015), and the one dealing with quantum computing (Bergholm et al. 2018; Anis et al. 2021; Broughton et al. 2020; Killoran et al. 2019; Baidu 2020). As the prototype becomes larger, the probability of introducing *bugs* in the code increases (Lipow 1982), possibly leading to erroneous results (Fidler et al. 2017; Botvinik-Nezer et al. 2020; Campos and Souto 2021). Well-organized code has been shown to facilitate code reuse and reproducibility (Trisovic et al. 2022; Mineault and Nozawa 2021). Minimizing the quantity of code needed to run an experiment has clear benefit in speeding up the research, reducing the time spent to learn, and put the software in a production environment.

We propose QuASK (Quantum Advantage Seeker with Kernel), a Python3 software framework unifying under a single interface all the features to run experiment with quantum kernels. QuASK can be run from the terminal using a single command line which specifies how to operate on the given data. Within the same command, the researcher can specify to analyze the data and subsequently generate graphics. QuASK can also be used as a library, to be integrated within an existing pipeline. Finally, the open-source nature of the framework allows the user to integrate further capabilities into the software, having them immediately available through the command line interface. QuASK is freely available at <https://github.com/CERN-IT-INNOVATION/QuASK> and the documentation is available at <https://quask.readthedocs.io/en/latest/index.html>.

2 Theoretical aspects of Quantum Kernels

A binary, symmetric function $K : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$ is a *kernel function* if positive definite (pd), i.e.

$$\sum_{i=1}^n \sum_{j=1}^n c_i c_j K(x_i, x_j) \geq 0 \quad (1)$$

for all $x_1, \dots, x_n \in \mathcal{X}$ given the real³ coefficients $c_1, \dots, c_n \in \mathbb{R}$. Supposing k continuous, we can associate a linear Hilbert-

¹ The linearity of the model means that the model can be expressed as linear transformation in the Hilbert space of the quantum system (e.g. matrix multiplication for finite-dimensional Hilbert spaces). However, the classifier is non-linear with respect to the space in which the original data lie, for the effect of the feature map.

² In some cases the quantum kernel values can concentrate around an average value, requiring a large number of shots to be estimated (Kübler et al. 2021; Thanasilp et al. 2022).

³ The above definition of real-valued kernel is true if we require also the kernel matrix to be symmetric.

Schmith integral operator

$$[Tg](x') = \int_{x \in \mathcal{X}} K(x', x)g(x)dx. \tag{2}$$

whose eigenfunctions $\{e_i\}_{i=1}^\infty$ form an orthonormal basis of square-integrable functions, and the sequence of corresponding eigenvalues $\{\lambda_i\}_{i=1}^\infty$ is non-negative. In such a case,

$$K(x, x') = \sum_{i=1}^\infty \lambda_i e_i(x)e_i(x') \tag{3}$$

For $x, x' \in \mathcal{X} \subseteq \mathbb{R}^d$, some examples of kernels are:

$$K_1(x, x') = x \cdot x' \quad \text{Linear} \tag{4}$$

$$K_p(x, x') = (x \cdot x' + b)^r \quad \text{Polynomial of degree } r \tag{5}$$

$$K_{\text{rbf}}(x, x') = \exp(-\alpha \|x - x'\|) \quad \text{RBF or Gaussian } (\alpha > 0) \tag{6}$$

$$K_{\text{eq}}(x, x') = 1 - \delta_{x,x'} \quad \text{Equality kernel} \tag{7}$$

while, for example, $|x - x'|$ is not a valid kernel due to the lack of positive definiteness. Important class of kernels are the *translation-invariant kernels* $K(x, x') = \Phi(x - x')$ given that \mathcal{X} is a vector space (the Gaussian kernel is an instance of such a family), and the *group kernel* $K(x, x') = \Phi(x^{-1}x')$ given \mathcal{X} has a group structure. Kernels can form other kernels: a non-negative linear combination of kernels, a product, and the limit of a kernel sequence (if exists) are kernels too. K_{eq} is a valid non-continuous kernel, which can be obtained as the limit for $n \rightarrow \infty$ of $\exp\{-n\|x - x'\|\}$. A larger list of kernels and their compositions can be found in (Rasmussen and Williams 2005; Duvenaud 2014).

Positive definite kernels can be thought as a generalization of the notion of inner product due to the strong relationship with the concept of *Reproducing Kernel Hilbert Space* (RKHS). A space $\mathcal{H} = \{f : \mathcal{X} \rightarrow \mathbb{R}\}$ of real-valued functions over \mathcal{X} is a RKHS if any linear functional $L_x : \mathcal{H} \rightarrow \mathbb{R}$, $L_x(f) = f(x)$ is bounded in \mathcal{H} (meaning that if two functions are close in terms of norm, then they are close also pointwise); or equivalently, for the Riesz representation theorem, it holds that $\forall x \in \mathcal{X}$ exists a unique $K_x \in \mathcal{H}$ such that $f(x) = L_x(f) = \langle f, K_x \rangle_{\mathcal{H}}$, $\forall f \in \mathcal{H}$. For every RKHS \mathcal{H} there is a unique K such that $K(x, x') = \langle K_x, K_{x'} \rangle$, namely K is a reproducing kernel for \mathcal{H} , and viceversa given a K positive definite kernel there is a unique Hilbert space of functions on \mathcal{X} for which K is a reproducing kernel (Aronszajn 1950). The mapping $\phi : \mathcal{X} \rightarrow \mathcal{H}$ encoding a data point within the RKHS is a *feature map*. Since

$$K(x, x') = \langle K_x, K_{x'} \rangle = \langle \phi(x), \phi(x') \rangle \tag{8}$$

we can interpret the application of K as calculating the inner product over a different vector space than the original point space \mathcal{X} .

In supervised learning applications is common to use a feature map ϕ to encode the data in higher dimensional (Hilbert) space to find a linear separation of the transformed data and by using the inverse map ϕ^{-1} we can recover a complex, nonlinear decision boundary in the original space. Due to the *representer theorem*, the linear pattern can be found independently of the dimensionality of \mathcal{H} : given the data points $\{(x_1, y_1), \dots, (x_m, y_m)\}$, the algorithms are fed with the kernel Gram matrix $K_{i,j} = [K(x_i, x_j)]$ of pairwise kernel similarities and no other information about the data is necessary for the classifier. Formally, the representer theorem asserts the linear function

$$\min_{f \in \mathcal{H}} L(f) + \lambda \|f\| \tag{9}$$

that minimizes the empirical risk is always in the form:

$$f(x) = \sum_{i=1}^m \alpha_i K(x, x_i). \tag{10}$$

The terms L is the loss function, e.g. the mean square error $L(f) = \frac{1}{m} \sum_{i=1}^m \|y - y'\|^2$. The term $\lambda \|f\|$, $\lambda > 0$ has regularization purposes, i.e. penalizes high norm solutions thus preferring smooth functions over non-smooth ones. The determination of the $\{\alpha_i\}$ values is a convex (efficient) optimization problem.

Kernel methods can be applied to supervised learning tasks using the kernel ridge regression algorithm (Murphy 2012), which is a straightforward generalization of linear regression, and the support vector machine (SVM) (Cortes and Vapnik 1995), which finds the linear classifier that maximizes the margin (i.e. the minimum distance between the data points and the boundary, on both sides). The SVM usually finds a sparse solution, i.e. a classifier whose output depends only on a few dataset items named support vectors. Kernel methods can be applied also to unsupervised learning tasks. Kernel PCA (Schölkopf et al. 1997) is the straightforward extension of the principal component analysis algorithm. It finds the components in the higher dimensional Hilbert space that have larger variance. Kernels can be applied to clustering techniques too, including the k-means algorithm (MacQueen 1967).

Kernel function can be parameterized, i.e. depending on one or more hyper-parameters that can be trained according to some loss function or chosen using a grid search. A different approach is the *multiple kernel learning* (Bach et al. 2004), which consists of defining multiple, fixed kernels and learning the most effective linear combination of such kernels.

2.1 Quantum Kernels implementation

Any parametric quantum circuit implementing the unitary transformation $U_\phi(\theta)$ acting on the Hilbert space \mathcal{H} of the n -qubits quantum system can be used to implement a feature map:

$$\begin{aligned} \phi : \mathcal{X} &\rightarrow \mathcal{H} \\ |0 \cdots 0\rangle &\mapsto U_\phi(x)|0 \cdots 0\rangle = |\phi(x)\rangle \end{aligned} \tag{11}$$

Such a feature map allows using the quantum space as an RKHS. In fact, we can obtain a kernel function sharing the same structure of Eq. 8 by encoding a pair of data points into quantum states and calculating the inner product between them:

$$\begin{aligned} K(x, x') &= \langle \phi(x) | \phi(x') \rangle \tag{12} \\ &= \text{prob}(\text{measurement of the state } U_\phi^\dagger(x')U_\phi(x)|0 \cdots 0\rangle \\ &\quad \text{using observable } \sigma_z^{\otimes n} \text{ collapse to eigenstate } |0 \cdots 0\rangle). \end{aligned} \tag{13}$$

Such a kernel can be concretely implemented using the overlap test circuit. The circuit structure is shown in Fig. 1a. We can equivalently use the SWAP test whose circuit structure is shown in Fig. 1b. Thus, we estimate the value of the kernel matrix $K_{i,j}$ by executing, for each pair of data points, U_ϕ multiple times (aka *shots*). This procedure consists in performing multiple measurements which force the quantum wavefunction to collapse, resulting in the fidelity measure between the two encoded data points. The kernel matrix can be finally fed to a kernel machine (e.g. SVM, Kernel PCA). Moreover, given the parametric quantum circuit for the feature map U_ϕ and a second quantum circuit W implementing the state $|w\rangle = W|0 \cdots 0\rangle$ corresponding to the linear weights, the function

$$f(x) = \langle \phi(x) | w \rangle \tag{14}$$

is a linear classifier.

Due to the large dimensionality of \mathcal{H} , exponentially in the number of qubits, the computation of the inner product may be affected by the *curse of dimensionality*: any two pairs of quantum states uniformly sampled in the Hilbert space have a high probability to be almost orthogonal (Ball et al. 1997). Each off-diagonal element of the Gram matrix vanishes with the increasing dimension of \mathcal{H} . If we think to perform the QKE on a current NISQ hardware this small value K_{ij} becomes indistinguishable from the inherently pervading noise making our classifier worthless. Such limitation requires a number of measurements to estimate a kernel function value that is polynomially in the dimensions of \mathcal{H} , thus exponentially in n . Therefore, we need to accurately design our unitary transformation in order to avoid losing quantum states within the Hilbert space.

We can design an effective quantum transformation, i.e. not affected by the curse of dimensionality, using several techniques.

The first approach is the use of parametric quantum circuits that we know analytically are restricted to a small subspace of $\mathcal{H}' \subset \mathcal{H}$, i.e. any parameter assignment x results in $U(x)|0 \cdots 0\rangle \in \mathcal{H}'$. The second approach is the use of a bandwidth coefficient (Canatar et al. 2022), i.e. a small scalar to be applied pointwise to the components of x , diminishing the range of each component. The third approach is to implement a projected or biased quantum kernel (Huang et al. 2021), which projects the quantum state to an approximate classical representation through an observable O (the choice of the observable is an educated guess). The quantum state lives in a large Hilbert space, but the observable O usually implies partial traces. The effect of a partial trace (present in O) over, e.g. the $k - th$ qubit, is to restrict the quantum space to some smaller representation, thus projecting it to the $k - th$ qubit subspace. A projected kernel function could take a gaussian form as follows:

$$k(x, x') = \exp(-\gamma \|\langle 0 | U_\phi(x) O U_\phi^\dagger(x) | 0 \rangle - \langle 0 | U_\phi(x') O U_\phi^\dagger(x') | 0 \rangle\|). \tag{15}$$

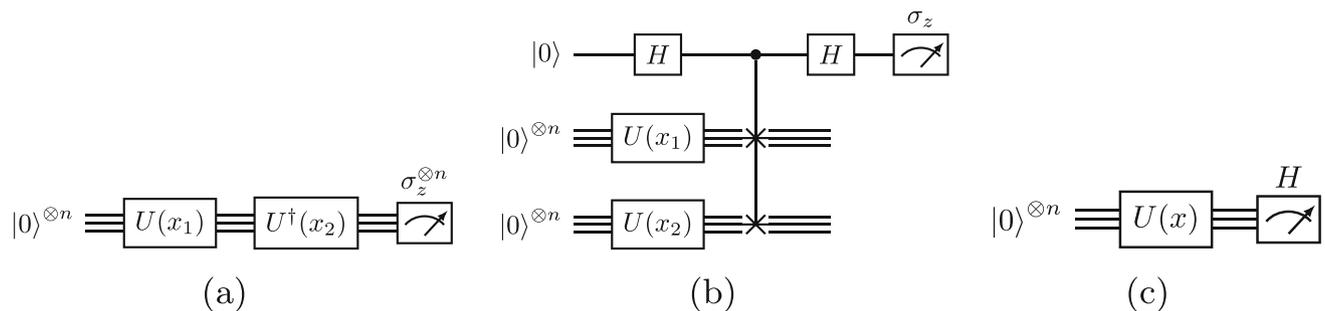


Fig. 1 Fig. 1(a)-(b) Fidelity test and SWAP test for Quantum Kernel Estimation where U is the feature map associated with the quantum kernel. Figure 1(c) Quantum circuit for the feature map associated with the projected kernel, the Hermitian observable H can be arbitrary

The transformation U_ϕ defining the quantum kernel influences dramatically the linear decision boundary to be found in the feature space through classical optimization. In fact, we can find an optimal⁴ form in an automatic fashion. The approach proposed by Glick et al. (2021) suggests the unitary transformation $U(x; \theta)$ should depend on both the data point features x and on some trainable parameters θ . The trainable parameters are then trained using stochastic gradient descent-based algorithm to minimize a loss function. Such an approach has been shown to be ineffective (Thanasilp et al. 2022). A different approach to optimize the parametric quantum circuit, choosing the basis gates at each point of the circuit as a combinatorial optimization algorithm (possibly a meta-heuristics) has been proposed by Incudini et al. (2022); Altares-López et al. (2021).

To evaluate the performance of a quantum kernel, a family of metrics has been introduced: geometric difference, approximate dimension, model complexity, and target-kernel alignment. The first three metrics constitute the central discussion of the Huang et al paper Huang et al. (2021). The last one has several implementations already. The ground-work can be found in Cristianini et al. (2001).

- The *geometric difference* compares classical and quantum kernel feature spaces evaluating the separation in performances of the two kernels. A large g compared with the \sqrt{N} indicates there is a deviation between the two kernel performances.
- The *approximate dimension* gives us an effective dimension of the quantum feature space generated by the encoding of the training samples. Indeed, this quantity helps us to understand the expressibility of the quantum kernels. If the d saturates with N it means the quantum states of the training data points are all orthogonal, otherwise a small value of d tells us the Hilbert space has not been fully exploited and the model has limited expressivity.
- The *model complexity* represents a final test where we find the complexity of a kernel including in the computation the labels of data. This metric derives from a prediction error generalization bound.
- The *target-kernel alignment*, as the model complexity, captures the relation present between a kernel and the relative target function, that is, the labels. The final objective of a kernel-based method is to approximate the label distribution with the data distribution in the feature space, and a margin tries to quantify this relation.

⁴ A form which guarantees good accuracy and generalization, or that satisfies some objective function as the metrics we refer to further in this section.

3 Quantum software frameworks

In recent years, a variety of software and programming languages have been developed to perform quantum computation. Most of the frameworks express quantum computation in terms of quantum circuits (Feynman 1985), which is the standard model *de facto*. They are usually able to apply a universal set of basis gates, decompose a unitary matrix to a quantum circuit, reverse a circuit, perform uncomputation (e.g. to restore the original value of an auxiliary qubit), and perform circuit transformation (e.g. replacing part of a circuit with another one). Such frameworks allow the simulation on the host computer, while others allow sending the quantum circuit to some remote quantum hardware to be executed. Some possible alternatives to the quantum circuit model are the quantum lambda calculus (Van Tonder 2004), the quantum Turing machine (Deutsch 1985), the adiabatic quantum model (Farhi et al. 2000), the measurement-based quantum computation (Raussendorf and Briegel 2001), the topological quantum computation (Kitaev 2003), and the ZX calculus (Coecke and Duncan 2011).

A comparison of frameworks using the quantum circuit model is shown in Table 1. Most frameworks allow to import and export of circuits in the OpenQASM format (Cross et al. 2017), an open-source specification for quantum circuits. This facilitates the porting of quantum software among the different platforms.

3.1 Quantum machine learning frameworks

PennyLane has been the first framework offering Quantum Machine Learning capabilities. They include the possibility to train a parametric quantum circuit, whose gradient can be calculated using the parameter-shift rule (Wierichs et al. 2022) or with finite difference method. It allows the integration of a quantum transformation as a layer in a neural network object defined in Keras (Chollet et al. 2015) or PyTorch (Paszke et al. 2019) libraries. PennyLane has also facilities to define a quantum kernel, whose fidelity circuit (Fig. 1a) is created automatically given the circuit for a quantum embedding. Strawberry Field proposes the same high-level capabilities for Continuous Variables formalism of quantum computing, like photonic Quantum Computing.

Qiskit Quantum Machine Learning has similar features, allowing us to embed quantum transformations within PyTorch networks and calculate kernel matrices. TensorFlow Quantum (Broughton et al. 2020) allows for rapid prototyping of hybrid-classical models due to its straightforward integration with the Machine Learning library TensorFlow (Abadi et al. 2015). Paddle Quantum allows for effortless application of QNNs to define LOCC (Local Operations and Classical Communication) protocols (Chitambar et al. 2014). Moreover, it allows the simulation of some quantum machine

Table 1 Comparison of relevant quantum computing frameworks

Vendor	Name	Ref	Language	S	E	QML
-	ProjectQ	Steiger et al. (2018)	Imperative	Yes	No	No
-	QCL	Ömer (2005)	Imperative	Yes	No	No
-	QiBO	Efthymiou et al. (2021)	Python	Yes	No	No
-	Quipper	Green et al. (2013)	Domain specific (Haskell)	Yes	No	No
-	Quirk	Gidney (2014)	Drag-and-drop	Yes	No	No
-	SiIQ	Bichsel et al. (2020)	Imperative	Yes	No	No
Amazon	Braket	Amazon Web Services (2020)	Python	Yes	Yes	No
Baidu	Paddle Quantum	Baidu (2020)	Python	Yes	No	Yes
Google	Cirq	Cirq Developers (2022)	Python	Yes	Yes	No
Google	TensorFlow Quantum	Broughton et al. (2020)	Python	Yes	Yes	Yes
IBM	Qiskit	Anis et al. (2021)	Python	Yes	Yes	Yes
Microsoft	Azure Quantum	Microsoft (2020)	Python, Q#	Yes	Yes	No
Microsoft	LIQUi)	Wecker and Svore (2014)	Domain specific (F#)	Yes	No	No
Quantinuum	t ket)	Sivarajah et al. (2020)	Python	Yes	Yes	No
Rigetti	Forest (pyQuil)	Rigetti (2019)	Python	Yes	Yes	No
Xanadu	PennyLane	Bergholm et al. (2018)	Python	Yes	Yes	Yes
Xanadu	Strawberry Field	Killoran et al. (2019)	Python	Yes	Yes	Yes

Legend: S: Simulation on CPU, E: Execution on quantum devices, QML: has quantum machine learning facilities. Domain-specific languages are embedded in the language denoted between parenthesis

learning algorithms defined in the measurement-based quantum computation formalism.

4 Proposed approach

As described in Section 3, many different quantum machine learning software exists, and most of them have few high-level algorithmic capabilities. However, there are several issues to address: many experiments require the interaction between different software platforms, e.g. Qiskit with PyTorch, requiring specific expert knowledge to be used. Furthermore, experiments need to be compared with theoretical results which are growing in the literature without, usually, a common implementation baseline.

Therefore, we have designed QuASK, a unifying, easy-to-use software framework that automates each phase of an experiment: the selection of the dataset, the preprocessing, the definition of the kernel, its implementation, and analysis. QuASK can be used both as a standalone executable through its command line interface and as a software library. The first approach performs the experiment without writing a single line of code. The second approach might be interesting if the researcher needs to use both existing code routines. After having accurately processed the data, and implemented them to compute classical and quantum Gram matrices we have a modest range of metrics (proposed in QuASK) to evaluate the obtained kernel methods.

4.1 Running experiments through a command line interface

We show how to use QuASK to perform an end-to-end experiment. Once installed, the software is run with `quask < command >`⁵. QuASK performs the sequence of operations illustrated in Fig. 2.

The experiment should start with the choice of a dataset. In such a case, QuASK offers several classical datasets both for regression and classification tasks. Moreover, some quantum datasets are available, i.e. datasets whose features has been encoded on a quantum system and modified by a unitary transformation, such as the one used in Huang et al. (2021). The output of the process is a pair of NumPy binary files representing the feature data and the corresponding labels.

```
> python3.9 -m quask get-dataset
The available datasets are:
0: iris - classification - classical
1: Fashion-MNIST - classification - classical
2: liver-disorders - regression - classical
3: delta_elevators - regression - classical
4: Q_Fashion-MNIST_2_E1 - regression - quantum
5: Q_Fashion-MNIST_4_E2 - regression - quantum
6: Q_Fashion-MNIST_8_E3 - regression - quantum
Which dataset will you generate?: 0
Where is the output folder?: .
Saved file ./iris_X.npy
Saved file ./iris_y.npy
```

The dataset, which can be obtained by using the QuASK command `get-command` or by using any dataset provided

⁵ It can be equivalently run with `python3.x -m quask < command >` where x is the Python version installed on the researcher's system.

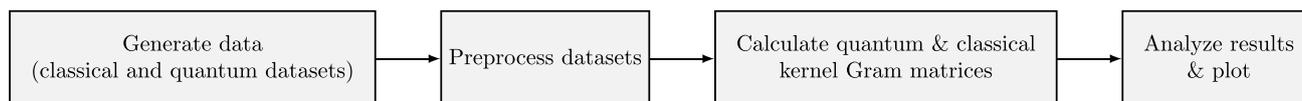


Fig. 2 Sequence of operations performed when analyzing a dataset using QuASK

by the user in NumPy format (a feature matrix X .`numpy`, $X \in \mathbb{R}^{d \times n}$, and a label vector y .`numpy`, $y \in \mathbb{R}^{1 \times n}$), can be preprocessed classically before being fed to the quantum machine learning algorithm. Several preprocessing techniques are available. Firstly, the researcher can vertically slice the dataset, keeping only a certain range of labels. Specifically, the software prompts the researcher to simplify the classification task by restricting it to binary classification. However, it is worth noting that most kernel-based predictors are able to handle both binary and multi-class classification problems. Secondly, the user can apply dimensionality reduction techniques. These are important especially in the NISQ setting due to the lack of resources. The techniques available are PCA for numerical data and FAMD for mixed numerical and categorical data.⁶ These choices are motivated by the fact that PCA is a widely used dimensionality reduction technique, while FAMD is a new method specifically tailored to handle categorical data. The user can extend QuASK to include further dimensionality reduction techniques. Thirdly, it is possible to fix the possible imbalance of the classes using random undersampling or random oversampling. When loading, the script already shows some statistics about the dataset, both for classification and regression tasks, which can guide the user through the preprocessing. The output of the process is the four files `X_train`, `y_train`, `X_test`, `y_test` which can now be fed to some kernel machine.

```

> python3.9 -m quask preprocess-dataset
Choose a random seed: 2323
Choose a random seed for splitting training and testing set: 3232
Where is the output folder?: .
Specify the path of classical feature data X (.numpy array): iris_X.npy
Loaded X having shape (150, 4) (150 samples of 4 features)
Specify the path of classical feature data Y (.numpy array): iris_y.npy
Loaded Y having shape (150,) (150 labels)
Labels have type int64 thus this is a classification problem
The dataset has 3 labels being [0, 1, 2] distributed as it follows:
  Class 0 is present 50 times (33.33 %)
  Class 1 is present 50 times (33.33 %)
  Class 2 is present 50 times (33.33 %)
Do you want to pick just the first two classes? [y/N]: y
Do you want to undersample the largest class? [y/N]: n
Do you want to oversample the smallest class? [y/N]: n
Do you want to apply preprocessing to the features? [y/N]: y
Do you want to apply PCA (numerical data only)? [y/N]: y
How many components you want?: 2
Do you want to apply FAMD (both numerical and categorical data)? [y/N]: n
Do you want to scale each field from 0 to 1 (MinMaxScaler)? [y/N]: y
Which percentage of data must be in the test set?: 0.5
Saved file ./X_train.npy
Saved file ./y_train.npy
Saved file ./X_test.npy
Saved file ./y_test.npy
  
```

At this point, the quantum kernel is built on the processed dataset. There are several available techniques the researcher can select from. The results of such an evaluation are the

kernel Gram matrices corresponding to the training and testing datasets.

```

> python3.9 -m apply-kernel
Where is X_train (numpy file)?: X_train.npy
Where is y_train (numpy file)?: y_train.npy
Where is X_test (numpy file)?: X_test.npy
Where is y_test (numpy file)?: y_test.npy
Do you want to apply a fixed kernel [Y] or a trainable one [N]? [y/N]: y
The available fixed kernels are:
0: linear_kernel
1: rbf_kernel
2: poly_kernel
3: zz_quantum_kernel
4: projected_zz_quantum_kernel
5: random_quantum_kernel
6: projected_random_quantum_kernel
Which kernel Gram matrix will you generate?: 0
Where is the output folder?: .
Saved file ./training_linear_kernel.npy
Saved file ./testing_linear_kernel.npy
  
```

The researcher can use *optimized quantum kernels*, i.e. quantum kernels whose circuits have been chosen after an optimization process. Such a process can be gradient-descent (ADAM optimizer) or gradient-free (grid search optimizer) based, in case we are optimizing the angles of the quantum operations, or combinatorial-optimization based, in case we are optimizing the generators of the quantum transformations. Although some quantum machine learning frameworks, such as PennyLane and Qiskit, already allows gradient-descent optimization of any circuit (including quantum kernels), no one offers the capabilities to adaptively choose the generators of the transformation through combinatorial optimization.

```

> python3.9 -m apply-kernel
Where is X_train (numpy file)?: X_train.npy
Where is y_train (numpy file)?: y_train.npy
Where is X_test (numpy file)?: X_test.npy
Where is y_test (numpy file)?: y_test.npy
Do you want to apply a fixed kernel [Y] or a trainable one [N]? [y/N]: n
Choose an embedding for your data (rx, ry, rz, zz): rx
Choose an embedding for your data (hardware_efficient, tfim, ltfim, zz_rx): tfim
Choose a number of layers: 1
Choose an optimizer (adam, grid): adam
Choose a reward metric to maximize (kernel-target-alignment, accuracy, geometric-difference, model-complexity): accuracy
Choose a random seed: 434343
WARNING:absl:No GPU/TPU found, falling back to CPU. (Set TF_CPP_MIN_LOG_LEVEL=0 and rerun for more info.)
.....
Where is the output folder?: .
Where is the output folder?: .
Saved file ./training_trainable_rx_tfim_1_adam_accuracy.npy
Saved file ./testing_trainable_rx_tfim_1_adam_accuracy.npy
  
```

Finally, the researcher can calculate the accuracy of the kernel model using the training and testing Gram matrices given as input. The output is a plot comparing the different kernels. For each kernel matrix, the user specifies the label that appears at the x-axis of the plot. If multiple instances are specified with the same label these are interpreted as i.i.d. random experiments and will contribute to the error bars. Multiple metrics are defined.

⁶ FAMD, Factor Analysis for Mixed Data is implemented using Prince library (available at <https://github.com/MaxHalford/prince>).

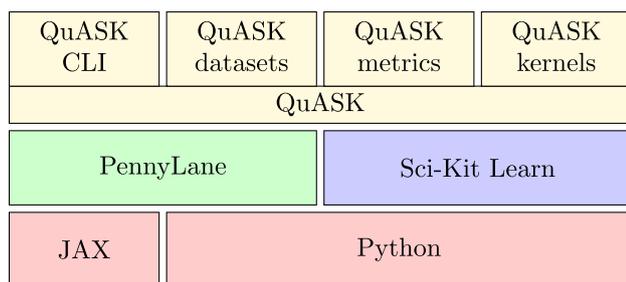


Fig. 3 Software stack describing the modularity of QuASK. QuASK is written on top of PennyLane for defining the quantum circuit software and Sci-Kit Learn. These frameworks, written in Python, allow accessing basic machine-learning routines

```
python3.9 -m quask plot-metric \
--metric accuracy \
--train-gram training_linear_kernel.npy --train-y Y_train.npy \
--test-gram testing_linear_kernel.npy --test-y Y_test.npy \
--label linear \
--train-gram training_rbf_01_kernel.npy --train-y Y_train.npy \
--test-gram testing_rbf_01_kernel.npy --test-y Y_test.npy \
--label rbf \
--train-gram training_rbf_001_kernel.npy --train-y Y_train.npy \
--test-gram testing_rbf_001_kernel.npy --test-y Y_test.npy \
--label rbf \
--train-gram training_poly_kernel.npy --train-y Y_train.npy \
--test-gram testing_poly_kernel.npy --test-y Y_test.npy \
--label poly \
--train-gram training_trainable_rx_tfim_1_adam_accuracy.npy \
--train-y Y_train.npy \
--test-gram testing_trainable_rx_tfim_1_adam_accuracy.npy --test-y Y_test.npy \
--label 'trainable quantum kernel'
```

4.2 Integrating QuASK in an existing code base

There might be cases in which the command line interface of QuASK cannot be straightforwardly used for a certain project. For example, the researcher might be forced to use a particular preprocessing technique, or analyze the result accordingly to a custom metric. Such cases can be addressed by integrating QuASK with the existing code base. In fact, QuASK provides a library of elements that can be integrated with other projects. The software is organized into several modules whose structure is shown in Fig. 3.

4.2.1 Download or generate datasets

The `quask.datasets` module facilitates the researcher in choosing a suitable dataset, providing some of the most popular datasets from OpenML platform⁷ and custom datasets generated also from quantum experiments (the latter allows us to reproduce results in Huang et al. (2021)).

```
from quask.datasets import load_who_life_expectancy_dataset

dataset = load_who_life_expectancy_dataset()
X = dataset['X']
y = dataset['y']
# use the dataset for any purpose
```

4.2.2 Evaluation metrics

The quality of a quantum kernel can be empirically tested through the performance of a kernel machine with respect to a certain dataset, by evaluating some metrics. The module `quask.metrics` contains the metrics to compare and evaluate the kernels, including the kernel polarity (i.e. Frobenius inner product between two Gram matrices), the target-kernel alignment (Cristianini et al. 2001), the training and testing accuracy of the Support Vector Machine with the precomputed kernel, the geometric difference (Huang et al. 2021) (which can be used to find a potential quantum advantage), and the model complexity (Huang et al. 2021).

```
from quask.metrics import calculate_approximate_dimension

kernel_gram_matrix = ...
d = calculate_approximate_dimension(kernel_gram_matrix)
```

4.2.3 Implement quantum kernels

The quantum kernel requires the definition of a feature map ϕ which is implemented using a parameterized unitary transformation $U_\phi(x)$. The Quantum Kernel Estimation algorithm (Havlíček et al. 2019) calculating the function $k(x, x') = \langle \phi(x) | \phi(x') \rangle$ is implemented through either the *overlap test* (Fig. 1a) or the *C-SWAP test* (Fig. 1b). The projected quantum kernel calculates classically the inner product between two feature vectors $\phi(x_1), \phi(x_2)$, each one being the output of the quantum feature map $\phi(x) = \langle 0 | U^\dagger(x) H U(x) | 0 \rangle$ (Fig. 1c). The feature map crosses the quantum space first through $U(x)$ and projects the data back in a classical representation when measuring with the Hermitian operator H . QuASK contains both some notable unitary transformations U from the literature and the code to use such unitary transformations as a kernel function through one of the three methods described above. The user can define their own unitary transformation and immediately get the kernel function. QuASK is agnostic with respect to the software framework used to define, simulate and execute the quantum circuits: we have implemented some unitary transformations in PennyLane. This allows also the use of the different functionalities offered by the different frameworks. For example, noiseless simulation with PennyLane can be speeded up using JAX.⁸ The open-source nature of QuASK allows for easy integration of other quantum computing frameworks in this platform.

The module `quask.kernels` collect all the quantum kernels defined within the platform. Most of the quantum kernels available are parametric quantum transformations in the form of Eq. 12. Such a module can be

⁷ OpenML is an open platform for sharing datasets, available at <https://www.openml.org/>

⁸ JAX is a high-performance linear algebra library.

straightforwardly extended to include user-specified quantum kernels. However, we can also have a more expressive quantum transformation, parameterized by both the user features and some trainable parameters, which can be adjusted using gradient-descent-based techniques to minimize some criteria. Such criteria can be one of the functions implemented in the `quask.metrics` module. The user can take advantage of the efficient `optax` optimization library. Moreover, we can use Structure Learning techniques (Incudini et al. 2022) to optimize the generators of the transformation using a combinatorial-optimization-based technique such as Simulated Annealing (Kirkpatrick et al. 1983) or Genetic Algorithms (Forrest 1996).

```
from quask.kernels import projected_zz_quantum_kernel
X = ... # feature matrix
kernel_gram_matrix = projected_zz_quantum_kernel(X, X, 0.1) # gamma = 0.1
```

4.3 Execution on real hardware

As the software is built on top of PennyLane, QuASK offers the same possibility of execution on real-world hardware. In particular, with the Qiskit-PennyLane plugin⁹ it is possible to run the quantum circuit on the IBM superconductor-based quantum hardware, and with the Braket-PennyLane plugin¹⁰ it is possible to exploit Rigetti, IonQ, and Oxford Quantum Circuit hardware. The execution on the NISQ hardware is noisy and the results may largely deviate from the simulated ones. The authors in Heyraud et al. (2022) have studied the effect of noise on quantum kernels.

5 Conclusions

We have introduced QuASK, a tool supporting researchers in creating powerful quantum kernels. The software takes care of the most time-consuming and error-prone aspects of the experimentation. It exploits theoretical metrics in QML, providing users with an environment to easily assess cases for potential quantum advantage. This package offers the exciting perspective of testing these metrics on real-world datasets. The QuASK project will be extended in future versions with a wider range of datasets and feature maps, both classical and quantum.

Acknowledgements FDM was supported by ERASMUS+ internship scholarship grant according to the EU regulation. MI is part of the Gruppo Nazionale Calcolo Scientifico of Istituto Nazionale di Alta Matematica Francesco Severi. MG is supported by CERN through the Quantum Technology Initiative.

⁹ <https://docs.pennylane.ai/projects/qiskit/en/latest>

¹⁰ <https://amazon-braket-pennylane-plugin-python.readthedocs.io/en/latest>.

Author Contributions All authors have contributed to the development of QuASK platform, and wrote the manuscript.

Data Availability The software is freely available at <https://github.com/CERN-IT-INNOVATION/QuASK>. The documentation is available at <https://quask.readthedocs.io/en/latest/index.html>.

Declarations

Ethics approval and consent to participate Not applicable.

Consent for publication Not applicable.

Conflict of interest Massimiliano Incudini and Davide Tezza are employees of Data Reply.

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

References

- Abadi M, Agarwal A, Barham P, Brevdo E, Chen Z, Citro C, Corrado GS, Davis A, Dean J, Devin M, Ghemawat S, Goodfellow I, Harp A, Irving G, Isard M, Jia Y, Jozefowicz R, Kaiser L, Kudlur M, Levenberg J, Mané D, Monga R, Moore S, Murray D, Olah C, Schuster M, Shlens J, Steiner B, Sutskever I, Talwar K, Tucker P, Vanhoucke V, Vasudevan V, Viégas F, Vinyals O, Warden P, Wattenberg M, Wicke M, Yu Y, Zheng X (2015) TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems. Software available from tensorflow.org. <https://www.tensorflow.org/>
- Abbas A, Sutter D, Zoufal C, Lucchi A, Figalli A, Woerner S (2021) The power of quantum neural networks. *Nat Comput Sci* 1(6):403–409
- Altares-López S, Ribeiro A, García-Ripoll JJ (2021) Automatic design of quantum feature maps. *Quantum Sci Technol* 6(4):045015
- Amazon Web Services (2020) Amazon Braket. <https://aws.amazon.com/braket/>
- Anis MS, Abby-Mitchell Abraham H et al (2021) Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2573505>
- Aronszajn N (1950) Theory of reproducing kernels. *Trans Am Math Soc* 68(3):337–404
- Arrasmith A, Cerezo M, Czarnik P, Cincio L, Coles PJ (2021) Effect of barren plateaus on gradient-free optimization. *Quantum* 5:558
- Bach, FR, Lanckriet GR, Jordan MI (2004) Multiple kernel learning, conic duality, and the SMO algorithm. In: Proceedings of the Twenty-first International Conference on Machine Learning. p 6
- Baidu (2020) Paddle Quantum. <https://github.com/PaddlePaddle/Quantum>
- Ball K et al (1997) An elementary introduction to modern convex geometry. *Flavors of Geometry* 31(1–58):26

- Bartkiewicz K, Gneiting C, Černoch A, Jiráková K, Lemr K, Nori F (2020) Experimental kernel-based quantum machine learning in finite feature space. *Sci Rep* 10(1):1–9
- Ben-Hur A, Ong CS, Sonnenburg S, Schölkopf B, Rätsch G (2008) Support vector machines and kernels for computational biology. *PLoS Comput Biol* 4(10):1000173
- Bergholm V, Izaac J, Schuld M, Gogolin C, Alam MS, Ahmed S, Arrazola JM, Blank C, Delgado A, Jahangiri S et al (2018) Pennylane: automatic differentiation of hybrid quantum-classical computations. [arXiv:1811.04968](https://arxiv.org/abs/1811.04968)
- Bharti K, Cervera-Lierta A, Kyaw TH, Haug T, Alperin-Lea S, Anand A, Degroote M, Heimonen H, Kottmann JS, Menke T, Mok W-K, Sim S, Kwek L-C, Aspuru-Guzik A (2022) Noisy intermediate-scale quantum algorithms. *Rev Mod Phys* 94:015004. <https://doi.org/10.1103/RevModPhys.94.015004>
- Bichsel B, Baader M, Gehr T, Vechev M (2020) SILQ: a high-level quantum language with safe uncomputation and intuitive semantics. In: Proceedings of the 41st ACM SIGPLAN Conference on Programming Language Design and Implementation. PLDI 2020, pp. 286–300. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3385412.3386007>
- Botvinik-Nezer R, Holzmeister F, Camerer CF, Dreber A, Huber J, Johannesson M, Kirchler M, Iwanir R, Mumford JA, Adcock RA et al (2020) Variability in the analysis of a single neuroimaging dataset by many teams. *Nature* 582(7810):84–88
- Broughton M, Verdon G, McCourt T, Martinez AJ, Yoo JH, Isakov SV, Massey P, Halavati R, Niu MY, Zlokapa A et al (2020) Tensorflow quantum: a software framework for quantum machine learning. [arXiv:2003.02989](https://arxiv.org/abs/2003.02989)
- Campos J, Souto A (2021) Qbugs: A collection of reproducible bugs in quantum algorithms and a supporting infrastructure to enable controlled quantum software testing and debugging experiments. In: 2021 IEEE/ACM 2nd International Workshop on Quantum Software Engineering (Q-SE). IEEE, pp 28–32
- Camps-Valls G (2006) Kernel Methods in Bioengineering, Signal and Image Processing. IGI Global, <https://doi.org/10.4018/978-1-59904-042-4>
- Canatar A, Peters E, Pehlevan C, Wild SM, Shaydulin R (2022) Bandwidth enables generalization in quantum kernel models. [arXiv:2206.06686](https://arxiv.org/abs/2206.06686)
- Chitambar E, Leung D, Mančinska L, Ozols M, Winter A (2014) Everything you always wanted to know about LOCC (but were afraid to ask). *Commun Math Phys* 328(1):303–326
- Chollet F et al (2015) Keras. GitHub. <https://github.com/fchollet/keras>
- Coecke B, Duncan R (2011) Interacting quantum observables: categorical algebra and diagrammatics. *New J Phys* 13(4):043016
- Cortes C, Vapnik V (1995) Support-vector networks. *Mach Learn* 20(3):273–297
- Cristianini N, Shawe-Taylor J, Elisseeff A, Kandola J (2001) On kernel-target alignment. *Advances in neural information processing systems* 14
- Cross AW, Bishop LS, Smolin JA, Gambetta JM (2017) Open quantum assembly language. [arXiv:1707.03429](https://arxiv.org/abs/1707.03429)
- Deutsch D (1985) Quantum theory, the church-turing principle and the universal quantum computer. *Proceedings of the Royal Society of London A Mathematical and Physical Sciences* 400(1818):97–117
- Developers C (2022) Cirq. <https://doi.org/10.5281/zenodo.6599601>
- Di Pierro A, Incudini M (2021) Quantum machine learning and fraud detection. *Protocols. Strands, and Logic*. Springer, Cham, Germany, pp 139–155
- Dumitrescu PT, Bohnet JG, Gaebler JP, Hankin A, Hayes D, Kumar A, Neyenhuis B, Vasseur R, Potter AC (2022) Dynamical topological phase realized in a trapped-ion quantum simulator. *Nature* 607(7919):463–467
- Duvenaud D (2014) Automatic model construction with Gaussian processes. PhD thesis, University of Cambridge
- Efthymiou S, Ramos-Calderer S, Bravo-Prieto C, Pérez-Salinas A, García-Martín D, García-Saez A, Latorre JI, Carrazza S (2021) Qibo: a framework for quantum simulation with hardware acceleration. *Quantum Science and Technology* 7(1):015018
- Farhi E, Goldstone J, Gutmann S, Sipser M (2000) Quantum computation by adiabatic evolution. [arXiv preprint quant-ph/0001106](https://arxiv.org/abs/quant-ph/0001106)
- Feynman RP (1985) Quantum mechanical computers. *Optics News* 11(2):11–20
- Fidler F, Chee YE, Wintle BC, Burgman MA, McCarthy MA, Gordon A (2017) Metaresearch for evaluating reproducibility in ecology and evolution. *Bioscience* 67(3):282–289
- Forrest S (1996) Genetic algorithms. *ACM Computing Surveys (CSUR)* 28(1):77–80
- Gidney C (2014) Quirk. GitHub. Available at: <https://github.com/Strilanc/Quirk>
- Glick JR, Gujarati TP, Corcoles AD, Kim Y, Kandala A, Gambetta JM, Temme K (2021) Covariant quantum kernels for data with group structure. [arXiv:2105.03406](https://arxiv.org/abs/2105.03406)
- Green AS, Lumsdaine PL, Ross NJ, Selinger P, Valiron B (2013) Quipper: a scalable quantum programming language. In: Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp 333–342
- Grossi M, Ibrahim N, Radescu V, Loredò R, Voigt K, Von Altröck C, Rudnik A (2022) Mixed quantum-classical method for fraud detection with quantum feature selection. [arXiv:2208.07963](https://arxiv.org/abs/2208.07963)
- Havlíček V, Córcoles AD, Temme K, Harrow AW, Kandala A, Chow JM, Gambetta JM (2019) Supervised learning with quantum-enhanced feature spaces. *Nature* 567(7747):209–212
- Heyraud V, Li Z, Denis Z, Le Boité A, Ciuti C (2022) Noisy quantum kernel machines. *Phys Rev A* 106(5):1
- Holmes Z, Arrasmith A, Yan B, Coles PJ, Albrecht A, Sornborger AT (2021) Barren plateaus preclude learning scramblers. *Phys Rev Lett* 126(19):190501
- Holmes Z, Sharma K, Cerezo M, Coles PJ (2022) Connecting ansatz expressibility to gradient magnitudes and barren plateaus. *PRX Quantum* 3(1):010313
- Huang H-Y, Broughton M, Mohseni M, Babbush R, Boixo S, Neven H, McClean JR (2021) Power of data in quantum machine learning. *Nature Communications* 12(1). <https://doi.org/10.1038/s41467-021-22539-9>
- Huang H-Y, Broughton M, Cotler J, Chen S, Li J, Mohseni M, Neven H, Babbush R, Kueng R, Preskill J et al (2022) Quantum advantage in learning from experiments. *Science* 376(6598):1182–1186
- Huang H-Y, Broughton M, Cotler J, Chen S, Li J, Mohseni M, Neven H, Babbush R, Kueng R, Preskill J, McClean JR (2022) Quantum advantage in learning from experiments. *Science* 376(6598):1182–1186. <https://doi.org/10.1126/science.abn7293>
- Incudini M, Martini F, Di Pierro A (2022) Structure learning of quantum embeddings. [arXiv:2209.11144](https://arxiv.org/abs/2209.11144)
- Killoran N, Izaac J, Quesada N, Bergholm V, Amy M, Weedbrook C (2019) Strawberry Fields: a software platform for photonic quantum computing. *Quantum* 3:129. <https://doi.org/10.22331/q-2019-03-11-129>. [arXiv:1804.03159](https://arxiv.org/abs/1804.03159)
- Kirkpatrick S, Gelatt CD Jr, Vecchi MP (1983) Optimization by simulated annealing. *Science* 220(4598):671–680
- Kitaev AY (2003) Fault-tolerant quantum computation by Anyons. *Ann Phys* 303(1):2–30
- Krunic Z, Flöthner FF, Seegan G, Earnest-Noble ND, Shehab O (2022) Quantum kernels for real-world predictions based on electronic

- health records. *IEEE Transactions on Quantum Engineering* 3:1–11
- Kübler J, Buchholz S, Schölkopf B (2021) The inductive bias of quantum kernels. *Advances in Neural Information Processing Systems* 34
- Kusumoto T, Mitarai K, Fujii K, Kitagawa M, Negoro M (2021) Experimental quantum kernel trick with nuclear spins in a solid. *npj Quantum Information* 7(1):1–7
- Kyriienko O, Magnusson EB (2022) Unsupervised quantum machine learning for fraud detection. [arXiv:2208.01203](https://arxiv.org/abs/2208.01203)
- Lipow M (1982) Number of faults per line of code. *IEEE Trans Software Eng* 4:437–439
- Liu N, Rebentrost P (2018) Quantum machine learning for quantum anomaly detection. *Phys Rev A* 97:042315. <https://doi.org/10.1103/PhysRevA.97.042315>
- Liu Y, Arunachalam S, Temme K (2021) A rigorous and robust quantum speed-up in supervised machine learning. *Nat Phys* 17(9):1013–1017
- MacQueen J (1967) Classification and analysis of multivariate observations. In: *Berkeley Symposium on Mathematical Statistics and Probability*. pp 281–297
- Madsen LS, Laudenbach F, Askarani MF, Rortais F, Vincent T, Bulmer JF, Miatto FM, Neuhaus L, Helt LG, Collins MJ et al (2022) Quantum computational advantage with a programmable photonic processor. *Nature* 606(7912):75–81
- Marrero CO, Kieferová M, Wiebe N (2021) Entanglement-induced barren plateaus. *PRX Quantum* 2(4):040316
- McClellan JR, Boixo S, Smelyanskiy VN, Babbush R, Neven H (2018) Barren plateaus in quantum neural network training landscapes. *Nat Commun* 9(1):1–6
- Mengoni R, Di Pierro A (2019) Kernel methods in quantum machine learning. *Quantum Mach Intell* 1(3):65–71
- Microsoft (2020) Azure Quantum SKD. <https://learn.microsoft.com/en-us/azure/quantum>
- Mineault P, Nozawa K (2021) patrickmineault/codebook: 1.0.0. Zenodo. <https://doi.org/10.5281/zenodo.5796873>
- Mitarai K, Negoro M, Kitagawa M, Fujii K (2018) Quantum circuit learning. *Phys Rev A* 98(3):032309
- Montanaro A (2016) Quantum algorithms: an overview. *npj Quantum Information* 2(1):1–8
- Murphy KP (2012) *Machine Learning: a Probabilistic Perspective*. MIT press, Cambridge, MA, USA
- Ömer B (2005) Classical concepts in quantum programming. *Int J Theor Phys* 44(7):943–955
- Paszke A, Gross S, Massa F, Lerer A, Bradbury J, Chanan G, Killeen T, Lin Z, Gimelshein N, Antiga L, Desmaison A, Kopf A, Yang E, DeVito Z, Raison M, Tejani A, Chilamkurthy S, Steiner B, Fang L, Bai J, Chintala S (2019) Pytorch: An imperative style, high-performance deep learning library. 8024–8035
- Pelofske E, Bärtschi A, Eidenbenz S (2022) Quantum volume in practice: what users can expect from NISQ devices. [arXiv:2203.03816](https://arxiv.org/abs/2203.03816)
- Pérez-Cruz F, Bousquet O (2004) Kernel methods and their potential use in signal processing. *IEEE Signal Process Mag* 21(3):57–65
- Peters E, Caldeira J, Ho A, Leichenauer S, Mohseni M, Neven H, Spentzouris P, Strain D, Perdue GN (2021) Machine learning of high dimensional data on a noisy quantum processor. *npj Quantum Information* 7(1):1–5
- Preskill J (2018) Quantum computing in the nisq era and beyond. *Quantum* 2:79
- Rasmussen CE, Williams CKI (2005) *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*. MIT Press, Cambridge, MA, USA
- Raussendorf R, Briegel HJ (2001) A one-way quantum computer. *Phys Rev Lett* 86(22):5188
- Rigetti (2019) Pyquil. <http://pyquil.readthedocs.io/en/latest>
- Rojo-Álvarez JL, Martínez-Ramón M, Muñoz-Mari J, Camps-Valls G (2018) *Digital Signal Processing with Kernel Methods*. John Wiley & Sons, New York, NY, USA
- Schölkopf B, Herbrich R, Smola AJ (2001) A generalized representer theorem. In: *International Conference on Computational Learning Theory*. Springer, pp 416–426
- Schölkopf B, Smola A, Müller K-R (1997) Kernel principal component analysis. In: *International Conference on Artificial Neural Networks*. Springer, pp 583–588
- Schuld M, Petruccione F (2021) *Machine learning with quantum computers*. Springer, Cham, Germany
- Schuld M, Killoran N (2019) Quantum machine learning in feature Hilbert spaces. *Phys Rev Lett* 122(4). <https://doi.org/10.1103/physrevlett.122.040504>
- Schuld M, Killoran N (2022) Is quantum advantage the right goal for quantum machine learning? *PRX Quantum* 3:030101. <https://doi.org/10.1103/PRXQuantum.3.030101>
- Sivarajah S, Dilkes S, Cowtan A, Simmons W, Edgington A, Duncan R (2020) t|ket>: a retargetable compiler for NISQ devices. *Quantum Science and Technology* 6(1):014003
- Steiger DS, Häner T, Troyer M (2018) Projectq: an open source software framework for quantum computing. *Quantum* 2:49
- Thanasilp S, Wang S, Cerezo M, Holmes Z (2022) Exponential concentration and untrainability in quantum kernel methods. [arXiv:2208.11060](https://arxiv.org/abs/2208.11060)
- Trisovic A, Lau MK, Pasquier T, Crosas M (2022) A large-scale study on research code quality and execution. *Scientific Data* 9(1):1–16
- Van Tonder A (2004) A lambda calculus for quantum computation. *SIAM J Comput* 33(5):1109–1135
- Wang G, Qi J (2014) Pet image reconstruction using kernel method. *IEEE Trans Med Imaging* 34(1):61–71
- Wang S, Fontana E, Cerezo M, Sharma K, Sone A, Cincio L, Coles PJ (2021) Noise-induced barren plateaus in variational quantum algorithms. *Nat Commun* 12(1):1–11
- Wang X, Du Y, Luo Y, Tao D (2021) Towards understanding the power of quantum kernels in the nisq era. *Quantum* 5:531
- Wecker D, Svore KM (2014) L|iquid>: a software design architecture and domain-specific language for quantum computing. [arXiv:1402.4467](https://arxiv.org/abs/1402.4467)
- Wierichs D, Izaac J, Wang C, Lin CY-Y (2022) General parameter-shift rules for quantum gradients. *Quantum* 6:677
- Yang M-H (2001) Face recognition using kernel methods. *Advances in neural information processing systems* 14