



A weighted quantum ensemble of homogeneous quantum classifiers

Emiliano Tolotti¹ · Enrico Blanzieri^{1,3} · Davide Pastorello^{2,3}

Received: 16 June 2025 / Accepted: 10 October 2025
© The Author(s) 2025

Abstract

Ensemble methods in machine learning aim to improve prediction accuracy by combining multiple models. This is achieved by ensuring diversity among predictors to capture different data aspects. Homogeneous ensembles use identical models, achieving diversity through different data subsets, and weighted-average ensembles assign higher influence to more accurate models through a weight learning procedure. We propose a method to achieve a weighted homogeneous quantum ensemble using quantum classifiers with indexing registers for data encoding. This approach leverages instance-based quantum classifiers, enabling feature and training point subsampling through superposition and controlled unitaries, and allowing for a quantum-parallel execution of diverse internal classifiers with different data compositions in superposition. The method integrates a learning process involving circuit execution and classical weight optimization, for a trained ensemble execution with weights encoded in the circuit at test-time. Empirical evaluation demonstrate the effectiveness of the proposed method, offering insights into its performance.

Keywords Quantum computing · Quantum machine learning · Ensemble methods · Quantum classifiers · Binary classification

1 Introduction

Ensemble methods are well-known in machine learning and they aim to enhance prediction accuracy by combining the outputs of multiple models (Polikar 2006). One important principle of ensemble methods is to ensure diversity among the predictors, allowing each of them to capture different aspects of the data. As a result, the combined models provide a more comprehensive understanding and better prediction capability. In the context of homogeneous ensembles, where the internal models are of the same kind, diversity can be introduced by using different sets of data for each classifier. The main methods include bagging (Breiman 1996), which considers subsets of training points sampled with replace-

ment, and attribute bagging (random subspaces) that uses random subsets of features (Ho 1998). Weighted-average ensembles are a subclass of ensemble methods where certain models are considered to be more accurate than others. These superior models receive higher weights through a weight learning procedure, giving them greater influence on the final prediction. They can be considered as stacking ensembles (Wolpert 1992), where the stacked classifier performs the weighted aggregation. By combining the predictions of individual models with varying weights, weighted-average ensembles leverage the strengths of each model to achieve superior performance.

The advantages of ensemble methods in the quantum setting have been presented in multiple works, where the proposed quantum-classical hybrid schemes show improvements in prediction performance (Zhang and Wang 2023; Qin et al. 2022; Incudini et al. 2023; West et al. 2023; Tolotti et al. 2024). The idea of quantum ensembles has been introduced in Schuld and Petruccione (2018); Abbas et al. (2020), where the authors propose a Bayesian Model Averaging method with non-trained classifiers weighted by accuracy. The authors in Macaluso et al. (2022, 2024) propose an ensemble strategy with the introduction of diversity by action of a controlled unitary execution acting on the data register. However, their approach involves a single exam-

✉ Emiliano Tolotti
emiliano.tolotti@unitn.it

¹ Department of Information Engineering and Computer Science, University of Trento, via Sommarive 9, 38123 Povo, Trento, Italy
² Department of Mathematics, Alma Mater Studiorum - Università di Bologna, Piazza di Porta San Donato 5, 40126 Bologna, Italy
³ Trento Institute for Fundamental Physics and Applications, via Sommarive 14, 38123 Povo, Trento, Italy

ple with qubit encoding, where individual data instances are encoded in different qubits, and the introduction of diversity by unitary application is performed entangling different training observations for each control state.

In contrast, we propose a method to achieve a weighted homogeneous quantum ensemble that employs a more efficient data encoding strategy based on instance-based quantum classifiers that rely on an indexing register, such as the quantum distance classifier (Schuld et al. 2017), the quantum cosine classifier (Pastorello and Blanzieri 2021), and a classification routine based on the SWAP test (Buhrman et al. 2001). These classifiers consider training instance vectors in superposition, requiring a different mechanism for introducing diversity. Additionally, rather than relying solely on diverse classifiers, which require independence to be advantageous, we propose a weighted aggregation of the internal models to increase accuracy.

Building on this, our method introduces diversity by allowing subsampling of both features and training points for these instance-based quantum classifiers. This is achieved through permutation of indexes and post-measurement state selection. By exploiting superposition and an additional control register that determines the size of the ensemble, the method enables quantum-parallel execution of internal classifiers. To facilitate effective weighting, our approach requires two distinct execution modalities, one for training and one for testing, differing in the measured qubits. We evaluated the performance of the proposed method using benchmarks on real-world datasets, which demonstrated improved accuracy over single classifiers, showcasing the effectiveness of the parallel weighted aggregation of diverse classifiers.

The remainder of this paper is organized as follows: Section 2 provides background information on the quantum classifiers used in our approach. Section 3 details our proposed method, including the quantum circuit design and the weight learning procedure. Section 4 presents experimental results demonstrating the effectiveness of our method. Section 5 discusses the implications of our findings and the role of the proposed ensemble. Finally, Section 6 draws some conclusions and suggests directions for future research.

2 Background on quantum classifiers

In this section we provide a brief overview of the quantum instance-based binary classifiers considered in this work. All the classifiers consider amplitude encoding of the data vectors, specifically they rely on a data register encoding the training set defined as follows:

$$|\psi_a\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle \sum_{j=0}^{M-1} x_i^j |j\rangle, \tag{1}$$

with

$$\sum_{j=0}^{M-1} \|x_i^j\|^2 = 1, \quad \forall i \in \{0, \dots, N-1\}, \tag{2}$$

where N and M indicate the number of data instances and features respectively, x_i^j indicates the feature j of the i -th data instance and $|i\rangle, |j\rangle$ are the basis states corresponding to the binary representation of i, j . Feature vectors are unit-norm normalized. All the classifiers consider a basis encoding for the binary labels $y_i \in \{-1, +1\}$, mapped to the domain $\{1, 0\}$ as

$$l_i = \frac{1 - y_i}{2}. \tag{3}$$

2.1 Quantum cosine classifier

The quantum cosine classifier (Pastorello and Blanzieri 2021) is a binary classification algorithm that is based on the cosine similarity between feature vectors. Its prediction function is defined as

$$y(x) = \text{sign} \left(\sum_{i=0}^{N-1} y_i \cos(x_i, x) \right), \tag{4}$$

with N being the number of training instances x_i with respective binary labels y_i . Its initial state is defined as


$$|\Psi\rangle = \frac{1}{\sqrt{2}}(|\psi_x\rangle|0\rangle + |\psi\rangle|1\rangle) \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_l \otimes \mathcal{H}_a, \tag{5}$$

with

$$|\psi_x\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle|x_i\rangle|l_i\rangle \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_l,$$

$$|\psi\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle|x\rangle|-\rangle \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_l.$$

The classifier requires the execution of a SWAP test (Buhrman et al. 2001) between an auxiliary qubit initialized in $|+\rangle$ and the ancilla a of the initial state $|\Psi\rangle_a$.



The probability of measuring 1 on the ancillary qubit controlling the SWAP test is

$$\mathbb{P}(1) = \frac{1}{4} \left(1 - \frac{2 \sum_{i=0}^{N-1} y_i \langle x | x_i \rangle}{\sqrt{2} \sum_{i=0}^{N-1} \|x\|^2 + \|x_i\|^2} \right). \tag{6}$$

With unit-norm data vectors, since

$$\cos(x, y) = \frac{x \cdot y}{\|x\| \|y\|}, \tag{7}$$

the prediction according to Eq. 4 is given by

$$y(x) = \text{sign}(1 - 4\mathbb{P}(1)). \tag{8}$$

2.2 Quantum distance classifier


The quantum distance classifier proposed in Schuld et al. (2017) is a binary classification algorithm based on the squared Euclidean distance between feature vectors. Its prediction function is defined as

$$y(x) = \text{sign} \left(\sum_{i=0}^{N-1} y_i \left(1 - \frac{1}{4} \|x - x_i\|^2 \right) \right), \tag{9}$$

where N is the number of training instances x_i with binary labels y_i . The initial state is defined as

$$|\Psi\rangle = \frac{1}{\sqrt{2N}} \sum_{i=0}^{N-1} |i\rangle(|0\rangle|x\rangle + |1\rangle|x_i\rangle)|l_i\rangle \in \mathcal{H}_n \otimes \mathcal{H}_a \otimes \mathcal{H}_m \otimes \mathcal{H}_l. \tag{10}$$

The circuit is composed by a Hadamard gate acting on the ancilla a and a conditional measurement of the label qubit l after measuring the ancilla.



Specifically, after a successful measurement of the ancilla qubit a in 0, the probability of measuring $k \in \{0, 1\}$ in l can be written as

$$\mathbb{P}(k|0) = \frac{1}{4Np_0} \sum_{i:l_i=k} \|x+x_i\|^2 = \frac{\sum_{i:l_i=k} \|x+x_i\|^2}{\sum_i \|x+x_i\|^2}, \tag{11}$$

where p_0 is the probability of measuring a in 0. With unit-norm data vectors, the following relationship holds:

$$\frac{1}{4N} \sum_i \|x+x_i\|^2 = \frac{1}{N} \sum_i \left(1 - \frac{1}{4} \|x-x_i\|^2 \right), \tag{12}$$

then, the predicted label according to Eq. 9 is given by

$$y(x) = \text{sign} \left(\mathbb{P}(0|0) - \frac{1}{2} \right). \tag{13}$$

2.3 Quantum SWAP test classifier


The authors in Blank et al. (2020) introduced a binary classification algorithm utilizing the SWAP test (Buhrman et al. 2001). This approach leverages a quantum kernel based on state fidelity. In our analysis, we consider a circuit that uses a single copy of the test data. However, this method can be extended to incorporate any number of copies, n , resulting in a kernel based on the n -th power of the fidelity. The prediction function is defined as

$$y(x) = \text{sign} \left(\sum_{i=0}^{N-1} y_i \cos^2(x_i, x) \right). \tag{14}$$

The initial state is defined as

$$|\Psi\rangle|x\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle|x_i\rangle|i_i\rangle|x\rangle \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_l \otimes \mathcal{H}_x, \tag{15}$$

where $|x\rangle$ indicates the state encoding the target vector x . The circuit applies the SWAP test between the qubits of the m and x registers, followed by a measurement of the control qubit of the SWAP a , and the label qubit l .



The joint probability of measuring q in the ancillary qubit and k in the label qubit can be written as

$$\mathbb{P}(q, k) = \frac{1}{2N} \sum_{i:l_i=k} \|x\|^2 \|x_i\|^2 + (-1)^q |\langle x | x_i \rangle|^2. \tag{16}$$

The classifier's output is computed from the expectation value of the observable corresponding to the measurement in the computational basis of the output qubits, as

$$\langle \sigma_z^a \otimes \sigma_z^l \rangle = \mathbb{P}(0, 0) - \mathbb{P}(0, 1) - \mathbb{P}(1, 0) + \mathbb{P}(1, 1)$$

$$= \frac{1}{N} \sum_{i=0}^{N-1} y_i |\langle x|x_i \rangle|^2. \tag{17}$$

Considering unit-norm data vectors, the prediction according to Eq. 14 is given by

$$y(x) = \text{sign}(\mathbb{P}(0, 0) - \mathbb{P}(0, 1) - \mathbb{P}(1, 0) + \mathbb{P}(1, 1)). \tag{18}$$

2.4 General formulation of the initial state

The classifiers described above share a common structure for the initial state, which we can express into a unified formulation. To provide a common notation for the different initializations, we express the initial states in a more general form that highlights the data encoding. We define

$$\begin{aligned} |\psi\rangle_{class} &= \sum_{i,j=0}^{2^n-1, 2^m-1} \alpha_{i,j} |i\rangle|j\rangle|\psi_i\rangle \\ &= \sum_{i,j=0}^{N-1, M-1} \alpha_{i,j} |i\rangle|j\rangle|\psi_i\rangle \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}, \end{aligned} \tag{19}$$

where we consider n and m qubits to encode the training samples and features respectively, so that $N \leq 2^n$ and $M \leq 2^m$. The coefficient $\alpha_{i,j}$ represents the amplitude of the j -th feature of the i -th data point, so $\alpha_{i,j} = 0, \forall i, j \notin [0, N - 1] \times [0, M - 1]$. In this expression, each classifier encodes a dataset of N training instances in superposition, with each instance represented by an indexing register (\mathcal{H}_n), identifying the training sample, and a feature register (\mathcal{H}_m) that encodes feature vectors through amplitude encoding. The auxiliary register (\mathcal{H}) stores class labels, ancillas, or other classifier-specific information.

We can express the initial state of the quantum classifiers in the unified form, by determining specific values for the amplitudes $\alpha_{i,j}$ and the auxiliary register state $|\psi_i\rangle$. For instance, the initial states of the quantum cosine and distance classifiers (referred to as *cdc*), described in Section 2.1 and Section 2.2 respectively, have a form of this type:

$$\begin{aligned} |\Psi^{cdc}\rangle &= \frac{1}{\sqrt{2N}} \sum_{i=0}^{N-1} |i\rangle(|x_i\rangle|\gamma\rangle|\phi_i\rangle \\ &+ |x\rangle|1 - \gamma\rangle|\theta_i\rangle) \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_a \otimes \mathcal{H}_c, \end{aligned} \tag{20}$$

where the input sample x is encoded in the same register of the training samples but entangled with different values $\gamma \in \{0, 1\}$ of the ancilla qubit a , and $|\phi_i\rangle, |\theta_i\rangle$ encode classifier-

specific information. To match the general form, the initial states of these classifiers can be rewritten as

$$\begin{aligned} |\Psi^{cdc}\rangle &= \frac{1}{\sqrt{2N}} \sum_{i,j=0}^{N-1, M-1} |i\rangle|j\rangle(x_i^j|\gamma\rangle|\phi_i\rangle \\ &+ x^j|1 - \gamma\rangle|\theta_i\rangle) = |\psi\rangle_{class}, \end{aligned} \tag{21}$$

where the amplitudes and the auxiliary register state are identified as

$$\begin{aligned} \alpha_{i,j}|\psi_i\rangle &= \frac{1}{\sqrt{2N}}(x_i^j|\gamma\rangle|\phi_i\rangle + x^j|1 - \gamma\rangle|\theta_i\rangle), \\ \forall i \in \{0, \dots, N - 1\}, \forall j \in \{0, \dots, M - 1\}. \end{aligned} \tag{22}$$

The SWAP test classifier, described in Section 2.3, encodes the test point x in an additional register (\mathcal{H}_x) instead, and the initial state can be generally written as follows:

$$|\Psi^{swap}\rangle = \frac{1}{\sqrt{N}} \sum_{i=0}^{N-1} |i\rangle|x_i\rangle|\psi_i\rangle|x\rangle \in \mathcal{H}_n \otimes \mathcal{H}_m \otimes \mathcal{H}_c \otimes \mathcal{H}_x. \tag{23}$$

This state can be recast into the unified form as

$$\begin{aligned} |\Psi^{swap}\rangle &= \frac{1}{\sqrt{N}} \sum_{i,j=0}^{N-1, M-1} x_i^j |i\rangle|j\rangle|\psi_i\rangle \left(\sum_{z=0}^{M-1} x_z |z\rangle \right) \\ &= |\psi\rangle_{class} \otimes |x\rangle, \end{aligned} \tag{24}$$

with

$$\begin{aligned} \alpha_{i,j}|\psi_i\rangle &= \frac{1}{\sqrt{N}}x_i^j|\psi_i\rangle, \\ \forall i \in \{0, \dots, N - 1\}, \forall j \in \{0, \dots, M - 1\}. \end{aligned} \tag{25}$$

3 Proposed method

In this section we outline the procedure and the details of the proposed quantum ensemble scheme.

3.1 Procedure

The procedure considers classifiers of the type reported in Section 2, whose initial states are expressed using the unified notation in Eqs. 21, 24. For presenting the procedure, we utilize the general form in Eq. 19, to which we add a control register and an additional ancilla qubit. The size (d) of the

control register determines the size of the ensemble, i.e. the number of diverse internal classifiers (2^d). The general initial state for the proposed ensemble, considering the control register, the initialized classifier and the additional ancilla qubit can be written as follows:

$$\begin{aligned}
 |\Phi_0\rangle &= |0\rangle_{ctrl}^{\otimes d} \otimes |\psi\rangle_{class} \otimes |0\rangle_{aux} \\
 &= |0\rangle^{\otimes d} \sum_{i,j=0}^{2^n-1, 2^m-1} \alpha_{i,j} |i\rangle|j\rangle|\psi_i\rangle|0\rangle. \tag{26}
 \end{aligned}$$

This general representation simplifies the description of the ensemble’s operations, and will be reconducted back to the classifier-specific formulations when it will be needed.

We discuss the steps of the procedure at test-time in Section 3.1.1, followed by the modifications utilized for training in Section 3.1.2.

3.1.1 Quantum ensemble circuit at test time

The test-time procedure performs the ensemble’s prediction for an input x , by executing the circuit exemplified in the diagram in Fig. 1. Suppose we have access to a positive unit-norm normalized weight vector w for aggregation of the internal classifiers of the ensemble. We encode the weights related to each internal classifier into the amplitudes of the control register by action of the unitary U_w :

$$|\Phi_1\rangle = U_w|\Phi_0\rangle = \sum_{c=0}^{2^d-1} \sqrt{w_c}|c\rangle \sum_{i,j=0}^{2^n-1, 2^m-1} \alpha_{i,j} |i\rangle|j\rangle|\psi_i\rangle|0\rangle, \tag{27}$$

where $\sqrt{w_c}$ indicates the amplitude related to the weight w_c .

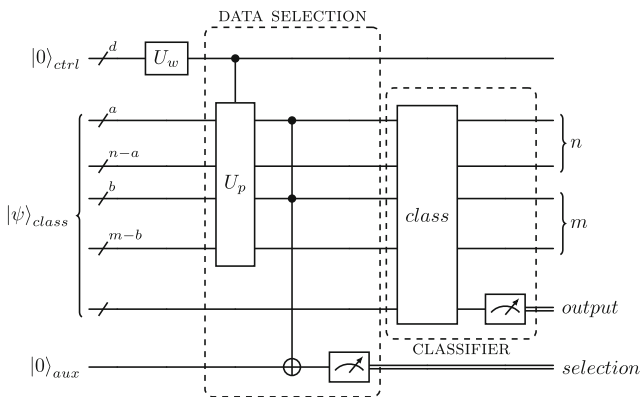


Fig. 1 Circuit diagram of the ensemble at test time. The control register is prepared with the weights encoded in the amplitudes. Then, the controlled permutation unitary is executed, acting on the data register of the classifier. The data register then partially controls (a qubits out of n and b qubits out of m) a NOT gate acting on the ancilla qubit, which is subsequently measured for data selection. The classifier is executed, and the related output qubit is measured. Specifically, the displayed circuit is designed to execute a quantum classifier of type cdc

In the next step, we apply controlled “permutation” unitary logic U_p , controlled by the control register and targeting the indexing register of training points and features. The purpose of the U_p unitary is to permute the elements of the training set in superposition, by entangling the data register with the control register. This can be achieved through bit-flip logic, as an application of CNOT and CSWAP gates. After the bit flipping, a different permutation of training samples and features is entangled with each basis state of the control register. The action of U_p can be written as follows:

$$|\Phi_2\rangle = U_p|\Phi_1\rangle = \sum_{c=0}^{2^d-1} \sqrt{w_c}|c\rangle \sum_{i,j=0}^{2^n-1, 2^m-1} \alpha_{i,j}^c |i\rangle|j\rangle|\psi_i\rangle|0\rangle, \tag{28}$$

where the amplitude $\alpha_{i,j}^c$ indicates a different feature amplitude depending on the control c for each index i, j .

After the permutation, a (multi-)controlled-NOT (MCX) gate is applied with control on partial index and feature registers, while acting on the ancillary qubit,

$$\begin{aligned}
 |\Phi_3\rangle &= U_{MCX}|\Phi_2\rangle \\
 &= \sum_{c=0}^{2^d-1} \sqrt{w_c}|c\rangle \left(\sum_{i,j \in sel} \alpha_{i,j}^c |i\rangle|j\rangle|\psi_i\rangle|0\rangle \right. \\
 &\quad \left. + \sum_{i,j \notin sel} \alpha_{i,j}^c |i\rangle|j\rangle|\psi_i\rangle|1\rangle \right). \tag{29}
 \end{aligned}$$

Referring to Fig. 1, assuming for example control by 2 qubits ($a = b = 1$), specifically the least significant in binary notation of the index and feature registers respectively, the action of the Toffoli gate (CCX) is

$$\begin{aligned}
 |\Phi_3\rangle &= U_{CCX}|\Phi_2\rangle \\
 &= \sum_{c=0}^{2^d-1} \sqrt{w_c}|c\rangle \sum_{i,j=0}^{2^n-1, 2^m-1} \left(\alpha_{2i,2j}^c |2i\rangle|2j\rangle|\psi_{2i}\rangle|0\rangle \right. \\
 &\quad + \alpha_{2i+1,2j}^c |2i+1\rangle|2j\rangle|\psi_{2i+1}\rangle|0\rangle \\
 &\quad + \alpha_{2i,2j+1}^c |2i\rangle|2j+1\rangle|\psi_{2i}\rangle|0\rangle \\
 &\quad \left. + \alpha_{2i+1,2j+1}^c |2i+1\rangle|2j+1\rangle|\psi_{2i+1}\rangle|1\rangle \right). \tag{30}
 \end{aligned}$$

The selection of training indexes and features is obtained by measuring the ancilla in the state $|\Phi_3\rangle$, obtaining a post-measurement state that reflects the corresponding state of the control qubits. Measuring 1 selects the training indexes and features whose binary representations have 1 in correspondence with the qubits that controlled the action on the ancilla, while measuring 0 selects the complementary set. The measurement probability is expected to be proportional to the size of the selected subset, assuming that the feature amplitudes are uniformly distributed. Measuring the ancilla collapses the state to a superposition of multiple diverse datasets, where each basis state of the control register is entangled with a

different encoded data composition. The post-measurement state results to be

$$|\Phi_4\rangle = \frac{1}{\sqrt{p_0}} \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle \sum_{i,j \in sel} \alpha_{i,j}^c |i\rangle |j\rangle |\psi_i\rangle |0\rangle, \quad (31)$$

where p_0 is the probability of measuring 0 in the ancilla qubit. Supposing we want to select the set where the ancilla is 0 in the $|\Phi'_3\rangle$ state ($p_0 \approx 0.75$), in the example above ($a = b = 1$) the resulting state after the data selection by measurement is as follows:

$$|\Phi'_4\rangle = \frac{1}{\sqrt{p_0}} \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle \sum_{i,j=0}^{2^{n-1}-1, 2^{m-1}-1} \left(\alpha_{2i,2j}^c |2i\rangle |2j\rangle |\psi_{2i}\rangle |0\rangle + \alpha_{2i+1,2j}^c |2i+1\rangle |2j\rangle |\psi_{2i+1}\rangle |0\rangle + \alpha_{2i,2j+1}^c |2i\rangle |2j+1\rangle |\psi_{2i}\rangle |0\rangle \right). \quad (32)$$

Note that in the case of cosine and distance classifiers (*cdc*) $|\Phi_4^{cdc}\rangle = |\Phi_4\rangle$, while in the case of a SWAP test (*swap*) classifier with separated test point register we have

$$|\Phi_4^{swap}\rangle = \frac{1}{\sqrt{p_0}} \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle \sum_{i,j \in sel} \alpha_{i,j}^c |i\rangle |j\rangle |x^{(c)}\rangle |\psi_i\rangle |0\rangle, \quad (33)$$

where the selection is performed on the test point register as well. Here, the permutation unitary U_p is applied also on $|x\rangle$, performing the same feature permutation of the training points also on the test point, for each basis state of the control register. In this case the circuit in Fig. 1, which is designed for the *cdc* classifier, should include also the test point register, and U_p should act on both the training and test point registers.

A reformulation of the preceding results can shed light on the nature of the elaboration. The equation Eq. 31 can be rewritten substituting back the original indexes and features, interpreting the selection as a different composition of the training set. In the case of the quantum cosine and distance classifiers Eq. 20, the state after the data selection can be written as

$$|\Phi_4^{cdc}\rangle = \frac{1}{\sqrt{p_0}} \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle \sum_{i \in sel} \delta_{c,i} |i\rangle (|x_i^{(c)}\rangle |\gamma\rangle |\psi_i\rangle + |x^{(c)}\rangle |1-\gamma\rangle |\phi_i\rangle) |0\rangle, \quad (34)$$

while in the case of the SWAP test classifier Eq. 23 it can be written as

$$|\Phi_4^{swap}\rangle = \frac{1}{\sqrt{p_0}} \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle \sum_{i \in sel} \delta_{c,i} |i\rangle |x_i^{(c)}\rangle |\psi_i\rangle |x^{(c)}\rangle |0\rangle. \quad (35)$$

In both cases, $x_i^{(c)}$ and $x^{(c)}$ are the selected training and test input points respectively, after applying the same permutation and selection of features, where $\delta_{c,i}$ represent normalization coefficients for each control basis state and training sample. It should be noted that the formulation resembles a superposition of the initial states of the classifiers, with the control register entangled with different compositions of the training and the test points.

The original circuit (denoted *class* in Fig. 1) of the classifier (*cdc* or *swap*) is then executed:

$$|\Phi_5\rangle = \sum_{c=0}^{2^d-1} \sqrt{w_c} |c\rangle |\phi_c\rangle |\varphi_c^{out}\rangle |0\rangle, \quad (36)$$

where $|\varphi_c^{out}\rangle$ denotes the output qubits state of each classifier, generally not normalized, and $|\phi_c\rangle$ indicates the rest of each classifier's state. Considering an observable O , the output of the ensemble classifier is given by the expectation value of O , obtained by collecting statistics of the measurement outcomes:

$$\mathbb{E}[O] = \frac{\sum_{c=0}^{2^d-1} w_c \langle \varphi_c^{out} | O | \varphi_c^{out} \rangle}{\sum_{c=0}^{2^d-1} w_c \langle \varphi_c^{out} | \varphi_c^{out} \rangle} = \sum_{c=0}^{2^d-1} w_c \langle \varphi_c^{out} | O | \varphi_c^{out} \rangle, \quad (37)$$

which resembles a linear combination of the outputs $\langle \varphi_c^{out} | O | \varphi_c^{out} \rangle$ of the internal classifiers, where the weights are represented by the unit-norm vector w .

As a special case, considering a single output qubit and projective measurement on the computational basis, $|\varphi_c^{out}\rangle = \beta_0^c |0\rangle + \beta_1^c |1\rangle$ and $O = |0\rangle\langle 0|$, Eq. 37 reduces to the following linear combination of outcome probabilities:

$$\mathbb{P}(0) = \frac{\sum_{c=0}^{2^d-1} w_c |\beta_0^c|^2}{\sum_{c=0}^{2^d-1} w_c (|\beta_0^c|^2 + |\beta_1^c|^2)} = \sum_{c=0}^{2^d-1} w_c |\beta_0^c|^2. \quad (38)$$

3.1.2 Quantum ensemble circuit at train time

The training procedure provides a way to obtain the weight vector w to be used at test time. To achieve this, it considers a set of samples as input and collects the outputs of each of the internal classifiers, which are used to determine the weights. The training circuit is similar to that during testing, with the corresponding diagram reported in Fig. 2, where the differences lie in the control register. During training, the classifiers are initialized on the training set and we consider uniform aggregation weights. The control register can be prepared in balanced superposition by applying Hadamard gates, replacing U_w with $H^{\otimes d}$ in Eq. 27. The same shuffling logic is then executed Eqs. 28, 29, and a measurement is taken to select the same data in superposition Eq. 31, so the ancilla output needs to match that at test time.

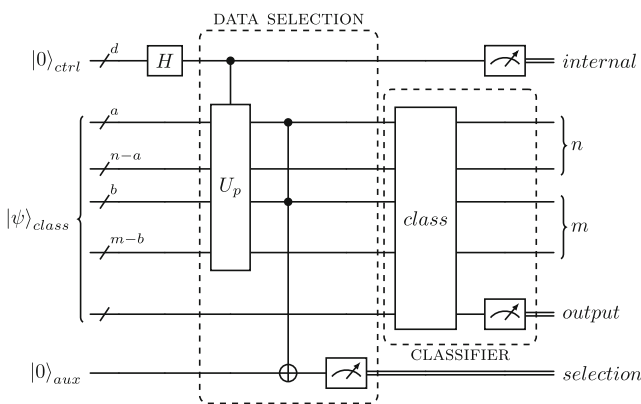


Fig. 2 Circuit diagram of the ensemble at training time. The control register is prepared in a balanced superposition using Hadamard gates. Then, the controlled permutation unitary is executed, acting on the data register of the classifier. The data register then partially controls a NOT gate acting on the ancilla qubit, which is subsequently measured for data selection. The classifier is executed, and finally, the control register is measured to get an internal classifier with the related output. Specifically, the displayed circuit is designed to execute a quantum classifier of type *cdc*

Upon a successful postselection (e.g., ancilla qubit in state $|0\rangle$), the classifier is executed:

$$|\Phi_5\rangle_{tr} = \frac{1}{2^{d/2}} \sum_{c=0}^{2^d-1} |c\rangle|\phi_c\rangle|\varphi_c^{out}\rangle|0\rangle. \tag{39}$$

The training phase requires measuring the control register to collect the predictions of the individual classifiers on a validation dataset. After measurement, the state collapses to that of a randomly chosen internal classifier c , determined by the measurement output of the control register:

$$|\Phi_6\rangle_{tr} = \frac{1}{\sqrt{p_c}} |c\rangle|\phi_c\rangle|\varphi_c^{out}\rangle|0\rangle, \tag{40}$$

where p_c is the probability of measuring the control register in the basis state corresponding to the classifier c , which is proportional to the norm of its output state:

$$p_c = \frac{\langle\varphi_c^{out}|\varphi_c^{out}\rangle}{2^d}. \tag{41}$$

Then, the observable O is measured on the output qubits of $|\varphi_c^{out}\rangle$. The prediction of the single internal classifier c is then:

$$\mathbb{E}[O]_c = \frac{\langle\varphi_c^{out}|O|\varphi_c^{out}\rangle}{\langle\varphi_c^{out}|\varphi_c^{out}\rangle}. \tag{42}$$

These values are estimated by executing the circuit multiple times, and are used to compute the aggregation weights w_c .

The training circuit execution is reported in Algorithm 1, which can be applied to predict the training set (or a different

Algorithm 1 Hybrid Training Procedure

```

Require: validation set  $D_{val}(x, y)$ , training circuit  $QC_{tr}$ 
Ensure: aggregation weights  $w$ 
1: for all  $x_i \in D_{val}(x, y)$  do
2:   for  $c \leftarrow 0$  to  $2^d - 1$  do
3:      $k_0[c] \leftarrow 0$  ▷ init 0 counts for classifier  $c$ 
4:      $k_1[c] \leftarrow 0$  ▷ init 1 counts for classifier  $c$ 
5:   end for
6:   for  $s \leftarrow 0$  to  $N_{shots} - 1$  do
7:      $c_i^s, o_i^s \leftarrow QC_{tr}(x_i)$  ▷ run the circuit for sample  $x_i$ 
8:     if  $o_i^s = 0$  then ▷ accumulate counts for classifier  $c_i^s$ 
9:        $k_0[c_i^s] \leftarrow k_0[c_i^s] + 1$ 
10:    else
11:       $k_1[c_i^s] \leftarrow k_1[c_i^s] + 1$ 
12:    end if
13:  end for
14:  for  $c \leftarrow 0$  to  $2^d - 1$  do
15:     $p[i, c] \leftarrow (k_0[c] + k_1[c]) / N_{shots}$  ▷ estimate  $p_c$  Eq. 41
16:     $p_0[i, c] \leftarrow k_0[c] / (k_0[c] + k_1[c])$  ▷ estimate  $\mathbb{P}(0)_c$  Eq. 42
17:  end for
18: end for
19:  $w \leftarrow weights\_learning(p, p_0)$  ▷ compute weights  $w$ 
20: return  $w$ 
    
```

validation). For simplicity, here we assume the output of the internal classifiers to be a single qubit, but the procedure can be generalized to multiple output qubits and observables. In this setting, $O = |0\rangle\langle 0|$ is the observable for measuring the output qubit. The expectation $\mathbb{E}[O]_c$ for the internal classifier c , corresponds to the probability $\mathbb{P}(0)_c$ of measuring 0 in the output qubit. The procedure estimates the values in Eqs. 41, 42, for each internal classifier c and for each validation sample x_i , that are used to compute the weights w for the aggregation of the internal classifiers. The algorithm iterates over the input samples x_i and runs the training circuit for a determined number of shots N_{shots} , related to the quality of the estimation. Each circuit execution outputs a random classifier index c_i^s and an output measurement o_i^s for that classifier and current data sample i . Here we assume the circuit to output a successful realization, but in practice the data selection circuit may fail, and the procedure should be repeated until a successful measurement is obtained. The results are then accumulated, for each internal classifier and validation sample, and the output values are estimated. Afterwards, the individual outputs are used to calculate the weights w for the classifiers to be used at test-time. These weights can be computed arbitrarily with classical or quantum procedures. A possibility is applying a weighting based on their performance on the training set, such as weighting by accuracy or by considering exponential weight decaying with errors. The chosen method involves optimization using logistic regression. The computed aggregation weights are then encoded in the amplitudes of the control register, taking into account positive weights. This enables us to consider an accuracy-weighted ensemble, having measured the performance in

predicting the training set, or to tune the weight vector with a model fit.

3.2 Data selection

The data selection procedure introduces diversity in the ensemble, by selecting the data subsets of each internal classifier. It is composed by a permutation of the dataset in superposition and a controlled action on an ancilla controlled by a few qubits of the data register with successive measurement of the ancilla. Regarding the permutation, it is performed by applying controlled operations acting on the data register controlled by the control register. This action, referred as U_p in Eq. 28, produces bit flips of the basis states of the data register changing its indexing differently for each basis state of the control register. The logic can be applied to both the index and feature registers of the classifier, achieving different ordering for the training instances and features producing 2^d potentially different amplitude permutations entangled with a control register of size d . For classifiers that consider the test point in a separate register in the initial state, as in Eq. 23, such as the the SWAP test classifier described in Section 2.3, the same controlled permutation logic used for the features needs to be applied also to the test register. In this way the different permutations of the features are consistent between training and test points.

To produce diversity for each classifier in superposition, we make the data register collapse to a subset of the indexes or features by measurement. We achieve that by a multi-controlled-NOT acting on an auxiliary qubit and controlled by a subset of the data register qubits (in Figs. 1 and 2, a and b qubits of the index and feature register respectively), as described in Eq. 29. The measurement of the auxiliary qubit collapses the states of the qubits that controlled the operation, effectively producing a subset of indexes of the data register, which will contain different data in superposition, as in Eq. 31. We considered a single ancillary qubit for this process, but an alternative approach is to consider multiple qubits to refine the data selection. For instance, one ancilla could be employed for selecting training indexes while the other focuses on the features, with each ancilla controlled by either the index or feature qubits, for defining more complex data selection patterns.

The permutation unitary and the ancilla control with measurement outcome are chosen, and deterministically decide the ensemble's internal classifiers data compositions and size with a given training set order. Their action can be arbitrary, with the requirement that it would produce different data sets which the ensemble can leverage from. For example the permutation can include multiple CSWAP and CNOT gates acting between different qubits, while the index selection

Algorithm 2 Data Selection Quantum Circuit

Require: control qubits, index qubits, feature qubits, ancilla
Ensure: circuit producing $2^{\text{control_qubits}}$ different data sets with $\approx \frac{3}{4}$ original size in superposition

```

1: for  $idx \leftarrow 0$  to  $control\_qubits - 1$  do
2:   if  $idx < \lfloor index\_qubits/2 \rfloor$  then
3:     CSWAP control(control[ $idx$ ])
         target(index[ $idx$ ], index[ $idx + \lfloor index\_qubits/2 \rfloor$ ])
4:   end if
5:   if  $idx < \lfloor feature\_qubits/2 \rfloor$  then
6:     CSWAP control(control[ $idx$ ])
         target(feature[ $idx$ ], feature[ $idx + \lfloor feature\_qubits/2 \rfloor$ ])
7:   end if
8:   if  $idx < index\_qubits$  then
9:     CNOT control(control[ $idx$ ]) target(index[ $idx$ ])
10:  end if
11:  if  $idx < feature\_qubits$  then
12:    CNOT control(control[ $idx$ ]) target(feature[ $idx$ ])
13:  end if
14: end for
15: for  $idx \leftarrow 0$  to  $control\_qubits - 1$  do
16:   for  $idx1 \leftarrow 0$  to  $control\_qubits - 1$  do
17:     if  $idx \neq idx1$  then
18:       CSWAP control(control[ $idx$ ])
         target(index[ $idx$ ], index[ $idx1$ ])
19:       CSWAP control(control[ $idx$ ])
         target(feature[ $idx$ ], feature[ $idx1$ ])
20:     end if
21:   end for
22: end for
23: CCX control(index[0], feature[0]) target(ancilla)
24:  $a \leftarrow$  measure ancilla
25: if  $a == 0$  then
26:   continue
27: else
28:   goto initial state preparation
29: end if

```

can be performed by considering different control qubits in the data register, and the measurement can be accepted if in 1, considering the subset of the indexes that activated the control, or in 0, considering the outer set. The success probability of this measurement, assuming the data is uniformly distributed, is proportional to the size of the subset we aim to select. Specifically, it approximately corresponds to the fraction of the subset relative to the total training set.

In this work we considered the data selection circuit described by the pseudocode in Algorithm 2, which employs CSWAP and CNOT operations for the controlled permutation, and for the selection the action of a CCNOT gate controlled by the first qubits of the index and feature registers and post-selection with 0 measurement. The circuit is designed to select a subset with a size of $\approx \frac{3}{4}$ of the original training set, with a success probability of $\approx \frac{3}{4}$. Specifically, the circuit discards the second half of the features the from the second half of the training points that are permuted differently in each internal classifier.

3.3 Weighted output combination

This section explains how the outputs of the internal classifiers are combined to produce the final prediction of the ensemble. At test time, the final ensemble prediction is the linear combination of the contributions β of each internal classifier, weighted according to the positive unit-norm weight vector w , encoded in the amplitudes of the control register, according to Eq. 37, which we can rewrite in vector notation as $w^T \beta$. To allow for both positive and negative weights, the vector w can be rescaled to be entirely positive by adding a bias b . In this case, the circuit needs to be executed twice for each input, with the rescaled positive weights $w + b$ and with uniform weights b (unit-norm normalized), and the final output is obtained by the difference of the runs, rescaled accordingly:

$$\frac{w^T \beta}{\|w\|} = \frac{1}{\|w\|} \left(\|w + b\| \frac{(w + b)^T \beta}{\|w + b\|} - \|b\| \frac{b^T \beta}{\|b\|} \right). \tag{43}$$

After obtaining the output probability of the classifier, by collecting statistics of the measurement outcome of the output qubit for multiple circuit repetitions, the output can be post-processed classically, obtaining the weighted ensemble output. In our case, the final decision is determined according to parameters computed with logistic regression. Specifically, as detailed in Sec. 3.1.2, during training the circuit is executed on a validation set to obtain the estimations of $p_{c,i}$ and $\mathbb{E}[O]_{c,i}$, for each internal classifier c and validation sample i (Eqs. 41 and 42), to be used for the parameters optimization. Then, we express the output of the ensemble classifier $\mathbb{E}[O]_i$ for each validation sample i , by means of the internal classifiers' outputs and weights w , by rewriting Eq. 37 as

$$\mathbb{E}[O]_i = \frac{\sum_{c=0}^{2^d-1} w_c \cdot p_{c,i} \cdot \mathbb{E}[O]_{c,i}}{\sum_{c=0}^{2^d-1} w_c \cdot p_{c,i}}. \tag{44}$$

The optimization minimizes the log-loss function:

$$\mathcal{L}(w, b, k) = -\frac{1}{|D_{val}|} \sum_{i=1}^{|D_{val}|} \left(y_i \log(\sigma(z_i)) + (1 - y_i) \log(1 - \sigma(z_i)) \right), \tag{45}$$

where $\sigma(z) = \frac{1}{1+e^{-z}}$ is the sigmoid function, $z_i = k\mathbb{E}[O]_i + b$, y_i is the label of the i -th sample, and $|D_{val}|$ is the number of validation samples. The positive weight vector w , a bias b , and a scaling factor k are optimized by bounded minimization of the log-loss using the L-BFGS-B algorithm (Zhu et al. 1997). Since the norm of the weight vector w does not influence the objective function in Eq. 45, the unnormalized w

obtained from the optimization is rescaled to unit norm for the encoding in the control register. By running the circuit with learned w as detailed in Sec. 3.1.1, Eq. 37 is estimated, and the prediction for an input x is then obtained as follows:

$$y(x) = \text{sign} \left(\sigma(z(x)) - \frac{1}{2} \right) = \text{sign}(\mathbb{E}[O] + b/k). \tag{46}$$

Considering a binary classification task, this formulation effectively performs a stacking of this logistic regression model on top of the internal classifiers, which act as a feature extractor layer. Considering the quantum classification algorithms and respective output probability formulations Eqs. 6, 11, 16, some of these methods require unit norm feature vectors for their prediction functions. However, in the data selection procedure, if we select a subset of features, the unit norm prerequisite is no longer valid. This is represented by the $\delta_{c,i}$ coefficients in Eqs. 34 and 35. Despite this, even if classifiers are not used as intended, utilizing a weight optimization procedure, such as the one we considered, allows the suboptimal internal classifiers' outputs (β) to serve as informative features leveraged through logistic regression.

4 Experiments

In this section we report empirical results of the performance of the proposed weighted ensemble.

4.1 Methodology

We implemented the scheme in Python with Qiskit (IBM 2024; Javadi-Abhari et al. 2024), and tested it on 11 real-world datasets from the UCI Machine Learning Library (Kelly et al. 2024), pre-processed to be suitable for a binary classification task. The implementation is available at <https://github.com/emiliantolo/quantum-ensemble-method>. The datasets are reported in Table 1.

The results have been computed with a Monte Carlo cross-validation technique (Xu and Liang 2001), with a 80%–20% training–validation dataset split and 10 independent runs. We employed the data with three different data normalization techniques. Specifically *none* refers to data as it is. Instead, *std* standardizes to mean 0 and standard deviation 1, then considering a training set $\{x_1, \dots, x_n\}$ the j -th feature of the i -th training instance is defined as

$$\text{std}(x_i^j) = \frac{x_i^j - \mu(x_1^j, \dots, x_n^j)}{\sigma(x_1^j, \dots, x_n^j)}, \tag{47}$$

where $\mu(x_1^j, \dots, x_n^j)$ and $\sigma(x_1^j, \dots, x_n^j)$ are mean and standard deviation of the j -th feature of the training set. With *min-max* each feature belongs to $[0, 1]$, then the normalization is

Table 1 Considered datasets

Name	# samples	# features	Class balance
iris_setosa_versicolor	100	4	balanced (50/50)
iris_setosa_virginica	100	4	balanced (50/50)
iris_versicolor_virginica	100	4	balanced (50/50)
vertebral_column_2C	310	6	unbalanced (100/210)
seeds_1_2	140	7	balanced (70/70)
ecoli_cp_im	220	7	unbalanced (77/143)
glasses_1_2	80	9	almost balanced (42/38)
breast_tissue_adi_fadmasgla	71	9	unbalanced (49/22)
breast_cancer	80	9	almost balanced (44/36)
accent_recognition_uk_us	80	12	unbalanced (63/17)
leaf_11_9	30	14	almost balanced (14/16)

defined as

$$\min\max(x_i^j) = \frac{x_i^j - \min_{i=1,\dots,n}(x_i^j)}{\max_{i=1,\dots,n}(x_i^j) - \min_{i=1,\dots,n}(x_i^j)}, \quad (48)$$

where $\min_{i=1,\dots,n}(x_i^j)$ and $\max_{i=1,\dots,n}(x_i^j)$ are the minimum and maximum features of the j -th feature in the training set. The normalization parameters have been computed on the training set samples for each dataset split. For the *min-max* normalization, test points exceeding the $[0, 1]$ bound are clipped.

Regarding the execution, *statevector simulation* refers to exact results computed based on the statevector, while *local simulation* indicates circuit execution in a noiseless environment with 8192 repetitions using the Aer simulator. It is worth noting that the number of shots includes both successful and unsuccessful data selections, hence the effective number of shots used for the estimation of the outputs is lower. The weighted ensemble adopts the data selection scheme described above and reported in Algorithm 2, thus the selection is performed by measuring the ancilla in 0 after the action of the first index and feature qubits, and the aggregation is by means of logistic regression which parameter optimization is performed using the SciPy package (Virtanen et al. 2020). We considered positive weights, constraining the optimizer with a positive bound. Moreover, we utilized the training set to train the weights at train-time.

4.2 Results

We considered the classifiers described above, namely the quantum cosine classifier, quantum distance classifier, quantum SWAP test classifier. We tested the single quantum classifiers using the full training set and simulated the execution of the ensemble's internal classifiers by applying the data selection procedure and feeding the resulting subset of the

original training set to the classifier. We evaluated ensembles with various sizes of the control register (where d indicates the number of control qubits).

Considering *statevector simulation*, we reported the results in Fig. 3. In the ideal case, we observed a performance degradation of the internal classifier compared to the single classifier on average, with some combinations of classifier and data normalization being more affected by data selection. Despite this, the weighted ensemble seems to effectively utilize these weaker internal classifiers, showing a noticeable average accuracy improvement across all considered data normalization methods and classifiers. Moreover, the mean accuracy improves slightly on average with increasing control register size. Additionally, we included XGBoost (Chen and Guestrin 2016), a parallel implementation of boosted decision trees, as a performance reference for classical ensembles. While it typically delivers the best performance, it is matched in accuracy by the weighted ensemble in certain conditions with these datasets.

Regarding the quantum cosine and distance classifiers, they exhibit similar behavior, highlighting their relationship, particularly considering their close prediction functions. This similarity is also reflected in their performance when used as feature extractors within the ensemble, where they show comparable performance across various data normalizations. Specifically, the best classification performance is achieved with *std* standardization. Conversely, *minmax* normalization results in some information loss. While it starts with acceptable accuracy, the ensemble's performance ultimately falls below that of non-normalized data (*none*).

Similarly, for the quantum SWAP test classifier, the ensemble with *none* achieves the best average performance, even though the single classifier performs best with *minmax*. On the other hand, *std* normalization, which centers features in zero, deteriorates the performance of the SWAP test classifiers since it is based on state fidelity, where squaring the values causes the sign to disappear.

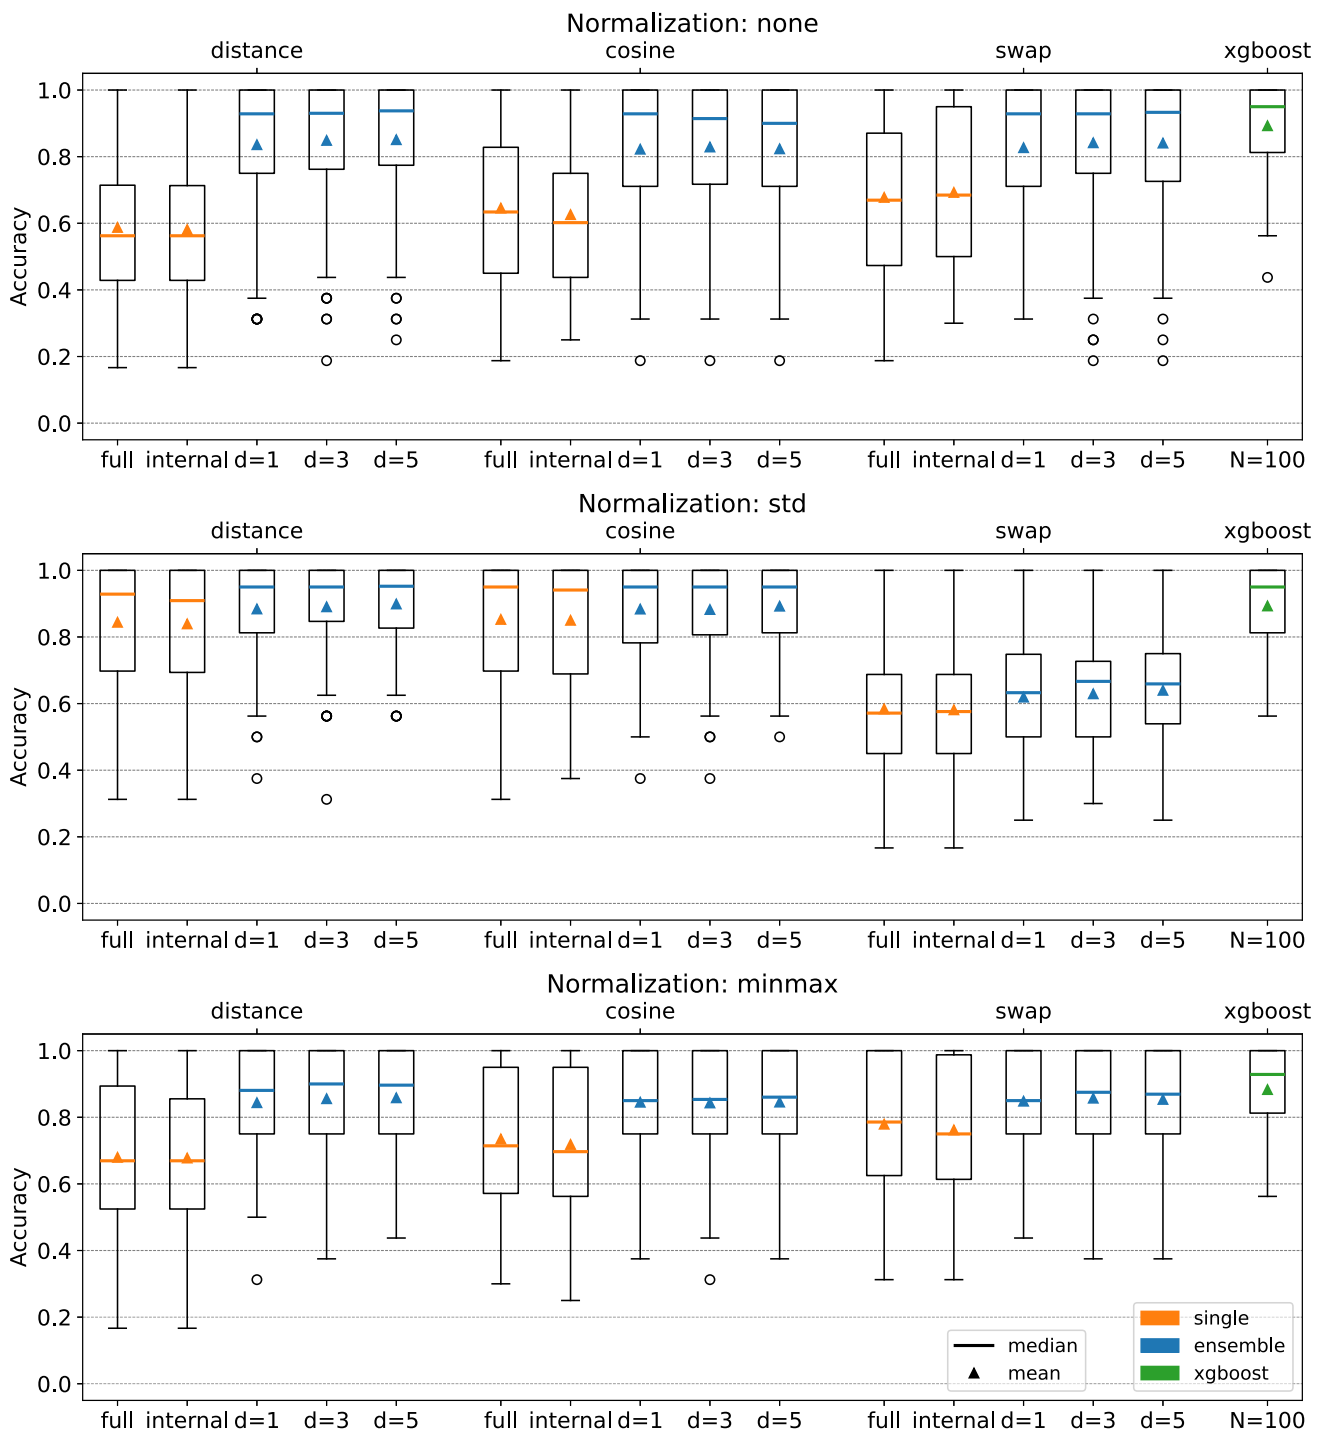


Fig. 3 Accuracy performance with *statevector simulation* of single classifiers (full), ensemble’s internal classifiers (internal), weighted ensemble and XGBoost in terms of aggregated results over 10 Monte Carlo runs for 11 real-world datasets. The considered internal models

are the quantum cosine, quantum distance and SWAP test classifiers. Multiple control register sizes (d) and normalization techniques are taken into account


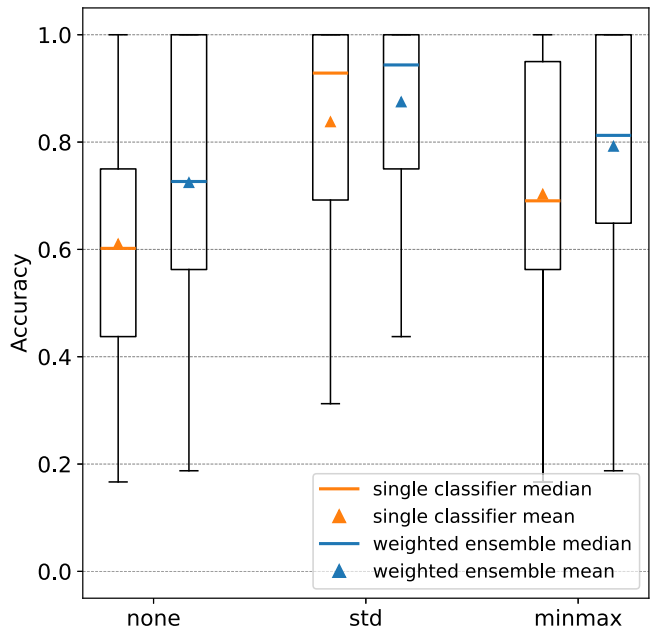


Fig. 4 Accuracy performance of *local simulation* with 8192 shots, considering the quantum distance classifier on 11 datasets over 10 Monte Carlo runs. The left panel shows the accuracy comparison between *local simulation* and *statevector simulation*, where points above the bisector



indicate accuracy degradation. The right panel presents the accuracy comparison of single and ensemble classifier ($d = 5$) with *local simulation* for the different normalization techniques

We conducted reduced-set testing using *local simulation* to demonstrate the impact of circuit repetitions. Due to computational constraints, we employed only the quantum distance classifier with the same data as the previous tests. The single distance classifier shows a minimal decrease in average accuracy with shot noise, as fluctuations of mispredicted points near the prediction threshold balance out those of correctly predicted points, implying the inaccuracy of the single classifier near the threshold. On the other hand, the weighted ensemble is more affected by measurement repetitions, likely due to error propagation in the optimization process. However, it still maintains an accuracy advantage over the single classifier. These results are displayed in Fig. 4.

5 Discussion

The proposed weighted quantum ensemble method demonstrates a significant performance improvement over individual quantum classifiers. This method allows for a theoretical exponential ensemble size with respect to the control register dimension.

In terms of qubits requirements, the ensembles introduces a modest increase relative to a single classifier. For a dataset of N instances and M features, the data register is composed of $n + m$ qubits with $n = \lceil \log_2 N \rceil$ and $m = \lceil \log_2 M \rceil$. A

single classifier then uses $n + m + c$ qubits, with c being the constant number of additional qubits required by the specific classifier. The number of qubits required by the ensemble is $d + n + m + c + 1$, as shown in Fig. 1, corresponding to $d + 1$ qubits above the single classifier, d for the control register and one ancilla for data selection.

In terms of computational complexity, the proposed weighted ensemble introduces a moderate overhead compared to a single classifier. At test time, the additional steps consist of preparing the control register (action of U_w in Fig. 1) and executing the data selection S (U_p and MCX in Fig. 1), both of which have a moderate impact on circuit cost in practice. In fact, the data encoding U_d , shared with the original classifier, is independent of U_w , and can therefore be executed in parallel prior to S , mitigating the effective overhead. Moreover, the data selection circuit S , detailed in Algorithm 2, scales polynomially with the number of control and data qubits ($T_S \sim \text{poly}(n, m, d)$), keeping the additional cost limited. The classifier execution C , with polynomial time complexity in the number of qubits T_C , remains unchanged. Specifically, if $T_{QC} = O(N_{\text{shots}} \cdot (T_{U_d} + T_C))$ is the time complexity of the single quantum classifier for N_{shots} repetitions of the circuit, then the complexity of the weighted ensemble at inference time is

$$T_{QE} = O \left(N_{\text{shots}} \cdot \frac{1}{F} \cdot (\max(T_{U_d}, T_{U_w}) + T_S + T_C) \right), \tag{49}$$

where $F < 1$ is the fraction of the training set selected by the data selection circuit. The term F reflects the fact that, on average, to obtain a useful computational path, the circuit must be executed $1/F$ times to select the correct data subset. The data selection we considered in Algorithm 2 corresponds to $F \approx 0.75$. Thus, the ensemble's inference complexity is comparable to that of the single classifier.

An important consideration is the cost of state preparation. In general, the exact amplitude encoding of an L -dimensional vector using $l = \lceil \log_2 L \rceil$ qubits requires a circuit depth of $O(L)$ (Möttönen et al. 2005; Iten et al. 2016). In this setting, the data encoding unitary U_d has a complexity of $T_{U_d} = O(N \cdot M)$, which also holds for the base classifiers (Pastorello and Blanzieri 2021; Schuld and Petruccione 2018; Blank et al. 2020). The weight encoding unitary U_w has cost $T_{U_w} = O(D)$, where $D = 2^d$ is the number of internal classifiers resulting from a d -qubit control register. In our scheme, $d \leq n + m$, so the control register is typically smaller than the data register, as in our experiments. Since U_d and U_w can be prepared in parallel, the overall preparation depth is determined by the larger of the two, usually U_d . In practice, T_{U_w} can often be neglected, and the additional overhead compared to the original classifiers is limited.

To mitigate the state preparation cost, several proposals aim to trade an increased number of ancillary qubits for reduced circuit depth. Specifically, for $L = 2^l$ amplitudes and l qubits we can achieve poly-logarithmic circuit depth $O(l^2)$ with $O(L^2)$ qubits (Zhang et al. 2021; Sun et al. 2023). Also, the authors of the SWAP-test classifier (Blank et al. 2020) indicate quantum forking as a possibility to prepare the initial state when each data sample is already available in separate registers, with linear routing depth in the number of data points. Moreover, approximation and compression strategies (e.g., Matrix Product State) have been proposed in the literature to mitigate state preparation cost (Malz et al. 2024), and their integration with the proposed ensemble remains an interesting direction for future work. While such approximations may introduce inaccuracies in the ensemble predictions, these effects could be mitigated by error-aware learning strategies, which might also address the impact of shot and hardware noise.

At training time, the ensemble circuit is executed to collect the outputs of the $D = 2^d$ internal classifiers for a validation set of size $|D_{val}|$. The time complexity for the training circuit execution is

$$T_{QE_{train}} = O\left(D \cdot |D_{val}| \cdot N_{shots} \cdot \frac{1}{F} \cdot (T_{U_d} + T_S + T_C)\right), \quad (50)$$

where $N_{shots} \cdot 1/F$ is the effective number of measurements for each internal classifier and validation sample. Notably, at this stage, there is no need to encode weights in the control register. The training procedure is hybrid as it requires an arbitrary classical optimization procedure to compute the

weights w . Since the optimization is performed on the outputs of the internal classifiers, the complexity of the optimization procedure is at least $O(D \cdot |D_{val}|)$. However, this is a one-time, offline procedure.

That said, the primary focus is not on computational complexity but rather on the performance enhancement achievable with the proposed scheme, which is evident even with limited control register sizes. Empirical results showed that the proposed weighted ensemble outperformed the single quantum classifier consistently, across various datasets and data normalization methods. This indicated that the ensemble successfully utilized the diversity of the weaker internal classifiers, introduced by the data selection. Moreover, the ensemble showed an increased robustness to data normalization, with instances where the ensemble with non-standardized data outperformed the ensemble with the standardization that showed the best result with the single classifier.

In the local simulation setting, where shot noise is present, the weighted ensemble was more sensitive to noise than the single classifier. This sensitivity is likely due to the optimization stage of the ensemble, which can be negatively affected by inaccuracies in individual outputs. Nonetheless, the ensemble showed acceptable resilience to noise, as it outperformed the accuracy of the single classifier.

Indeed, for linear classifiers, such as distance and cosine quantum classifiers, the weighted aggregation can be interpreted as a modified dataset with multiplicative factors on training points and features. This allows for an alternative execution of the weighted ensemble at test-time by using a single quantum classifier on the modified dataset, avoiding the data selection circuit logic but requiring a change in the data encoding procedure of the initial state.

For non-linear classifiers, such as those using the SWAP test, the weighted ensemble can be interpreted as a two-layer neural network (NN) with some missing connections in the first layer and the second layer's weights trained by an optimization procedure. A similar NN structure based on the SWAP test has been proposed in Pastorello and Blanzieri (2024), which considers a classical combination of quantum modules, whereas the proposed method has a quantum aggregation of SWAP test based classifiers. To assess the non-linear prediction capabilities and correctness of our approach, we tested it on a synthetic 2-dimensional XOR dataset. As expected, the linear classifiers and related ensembles struggled, while the ensemble using the non-linear SWAP test achieved nearly perfect accuracy.

Additionally, a majority voting aggregation strategy can be implemented by sampling the circuit defined for the training procedure. This allows the individual outputs of the internal classifiers to be estimated and then combined in a bagging fashion via majority voting. Sampling with replacement-like data sampling may be advantageous and

can be simulated by using a larger training dataset, employing multiple copies of the original training data, and allowing for training point repetitions in each subset up to the number of original training set copies.

6 Conclusion

In this work, we proposed a weighted homogeneous quantum ensemble method that leverages instance-based quantum classifiers to improve classification performance. The method introduces diversity through training point and feature subsampling in superposition, allowing for a quantum-parallel execution of the diverse internal classifiers, which outputs are combined with weights. The weight learning procedure is hybrid, where the quantum circuit is executed to collect the outputs of the internal classifiers, and the weights are trained classically. The deployment is quantum, encoding the weights in the circuit at test-time. Importantly, the ensemble introduces a modest computational overhead, as the dominant complexity remains that of the underlying original quantum classifier.

The empirical evaluation on real-world datasets demonstrated the effectiveness of the proposed method. The results showed improved accuracy over the single quantum classifier, both in the exact case by considering results computed with expectation values, and with shot noise considering circuit measurement repetitions. We compared the accuracy of the quantum ensemble with XGBoost and showed comparable performance with some internal classifiers and data normalizations.

Moreover, let us remark that, while variational quantum circuits have become central in current QML research, non-variational approaches like the proposed one based on ensembles of similarity-based classifiers, avoiding gradient-based optimization for circuit training and related issues (e.g. barren plateaus), align naturally with the NISQ constraints. In this way, in our opinion, quantum ensembles of quantum (non-variational) classifiers represent an advantageous solution in QML.

Future directions include exploring different diversity introduction mechanisms not based on data subsetting, such as parametric controlled rotation operations. Moreover, we aim at investigating approximated state preparation techniques as a way to reduce circuit depth, and their impact on the ensemble prediction performance. Overall, the proposed quantum weighted ensemble method offers a robust and flexible framework for improving classification accuracy in quantum machine learning.

Acknowledgements D.P. is a member of the “Gruppo Nazionale per la Fisica Matematica (GNFM)” of the “Istituto Nazionale di Alta Matematica “Francesco Severi” (INdAM)”.

Author Contributions Conceptualization: E.T., E.B., D.P.; Methodology: E.T., E.B., D.P.; Formal analysis and investigation: E.T., E.B., D.P.; Code implementation and experiments: E.T.; Writing - original draft preparation: E.T., E.B., D.P.; Supervision: E.B., D.P.;

Funding Open access funding provided by Università degli Studi di Trento within the CRUI-CARE Agreement. E.T. was supported by the MUR National Recovery and Resilience Plan (PNRR) M4C1I4.1, funded by the European Union under NextGenerationEU. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or The European Research Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. D.P. was supported by project SERICS (PE00000014) under the MUR National Recovery and Resilience Plan funded by the European Union – NextGenerationEU.

Data Availability Code and datasets are available at <https://github.com/emiliantolo/quantum-ensemble-method>.

Declarations

Competing interests The authors declare no competing interests.

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

- Abbas A, Schuld M, Petruccione F (2020) On quantum ensembles of quantum classifiers. *Quantum Mach Intell* 2(1):6. <https://doi.org/10.1007/s42484-020-00018-6>
- Blank C, Park DK, Rhee J-KK, Petruccione F (2020) Quantum classifier with tailored quantum kernel. *npj Quantum Inf* 6(1):1–7. <https://doi.org/10.1038/s41534-020-0272-6>
- Breiman L (1996) Bagging predictors. *Mach Learn* 24(2):123–140. <https://doi.org/10.1007/BF00058655>
- Buhrman H, Cleve R, Watrous J, de Wolf R (2001) Quantum Fingerprinting. *Phys Rev Lett* 87(16):167902. <https://doi.org/10.1103/PhysRevLett.87.167902>
- Chen T, Guestrin C (2016) XGBoost: A Scalable Tree Boosting System. <https://doi.org/10.1145/2939672.2939785>
- Ho TK (1998) The random subspace method for constructing decision forests. *IEEE Trans Pattern Anal Mach Intell* 20(8):832–844. <https://doi.org/10.1109/34.709601>
- IBM (2024) Qiskit | IBM Quantum Computing. <https://www.ibm.com/quantum/qiskit>
- Incudini M, Grossi M, Ceschini A, Mandarino A, Panella M, Vallecorsa S, Windridge D (2023) Resource saving via ensemble techniques for quantum neural networks. *Quantum Mach Intell* 5(2):39. <https://doi.org/10.1007/s42484-023-00126-z>

- Iten R, Colbeck R, Kukuljan I, Home J, Christandl M (2016) Quantum circuits for isometries. *Phys Rev A* 93(3):032318. <https://doi.org/10.1103/PhysRevA.93.032318>
- Javadi-Abhari A, Treinish M, Krsulich K, Wood CJ, Lishman J, Gacon J, Martiel S, Nation PD, Bishop LS, Cross AW, Johnson BR, Gambetta JM (2024) Quantum Computing with Qiskit. <https://doi.org/10.48550/arXiv.2405.08810>
- Kelly M, Longjohn R, Nottingham K (2024) Home - UCI Machine Learning Repository. <https://archive.ics.uci.edu/>
- Macaluso A, Clissa L, Lodi S, Sartori C (2024) An efficient quantum algorithm for ensemble classification using bagging. *IET Quantum Commun* 5(3):253–268. <https://doi.org/10.1049/qt2.12087>
- Macaluso A, Clissa L, Lodi S, Sartori C (2022) Quantum Ensemble for Classification. <https://doi.org/10.48550/arXiv.2007.01028>
- Malz D, Styliaris G, Wei Z-Y, Cirac JI (2024) Preparation of Matrix Product States with Log-Depth Quantum Circuits. *Phys Rev Lett* 132(4):040404. <https://doi.org/10.1103/PhysRevLett.132.040404>
- Möttönen M, Vartiainen JJ, Bergholm V, Salomaa MM (2005) Transformation of quantum states using uniformly controlled rotations. *Quantum Info Comput* 5(6):467–473
- Pastorello D, Blanzieri E (2021) A Quantum Binary Classifier based on Cosine Similarity. In: 2021 IEEE International conference on quantum computing and engineering (QCE), pp 477–478. <https://doi.org/10.1109/QCE52317.2021.00086>
- Pastorello D, Blanzieri E (2024) Scalable quantum neural networks by few quantum resources. *Int J Quantum Inf* 2450018. <https://doi.org/10.1142/S0219749924500187>
- Polikar R (2006) Ensemble based systems in decision making. *IEEE Circuits Syst Mag* 6(3):21–45. <https://doi.org/10.1109/MCAS.2006.1688199>
- Qin R, Liang Z, Cheng J, Kogge P, Shi Y (2022) Improving Quantum Classifier Performance in NISQ Computers by Voting Strategy from Ensemble Learning. <https://doi.org/10.48550/arXiv.2210.01656>
- Schuld M, Petruccione F (2018) Quantum ensembles of quantum classifiers. *Sci Reports* 8(1):2772. <https://doi.org/10.1038/s41598-018-20403-3>
- Schuld M, Fingerhuth M, Petruccione F (2017) Implementing a distance-based classifier with a quantum interference circuit. *Europhys Lett* 119(6):60002. <https://doi.org/10.1209/0295-5075/119/60002>
- Sun X, Tian G, Yang S, Yuan P, Zhang S (2023) Asymptotically Optimal Circuit Depth for Quantum State Preparation and General Unitary Synthesis. <https://doi.org/10.48550/arXiv.2108.06150>
- Tolotti E, Zardini E, Blanzieri E, Pastorello D (2024) Ensembles of quantum classifiers. *Quantum Inf Comput* 24(3&4):181–209. <https://doi.org/10.26421/QIC24.3-4-1>
- Virtanen P, Gommers R, Oliphant TE, Haberland M, Reddy T, Cournapeau D, Burovski E, Peterson P, Weckesser W, Bright J, van der Walt SJ, Brett M, Wilson J, Millman KJ, Mayorov N, Nelson ARJ, Jones E, Kern R, Larson E, Carey CJ, Polat İ, Feng Y, Moore E.W, VanderPlas J, Laxalde D, Perktold J, Cimrman R, Henriksen I, Quintero EA, Harris CR, Archibald AM, Ribeiro AH, Pedregosa F, van Mulbregt P (2020) SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nat Methods* 17(3):261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- West MT, Seviar M, Usman M (2023) Boosted Ensembles of Qubit and Continuous Variable Quantum Support Vector Machines for B Meson Flavor Tagging. *Adv Quantum Technol* 6(10):2300130. <https://doi.org/10.1002/qute.202300130>
- Wolpert DH (1992) Stacked generalization. *Neural Netw* 5(2):241–259. [https://doi.org/10.1016/S0893-6080\(05\)80023-1](https://doi.org/10.1016/S0893-6080(05)80023-1)
- Xu Q-S, Liang Y-Z (2001) Monte Carlo cross validation. *Chemometrics Intell Laboratory Syst* 56(1):1–11. [https://doi.org/10.1016/S0169-7439\(00\)00122-2](https://doi.org/10.1016/S0169-7439(00)00122-2)
- Zhang X-Y, Wang M-M (2023) An efficient combination strategy for hybrid quantum ensemble classifier. *Int J Quantum Inf* 21(06):2350027. <https://doi.org/10.1142/S0219749923500272>
- Zhang X-M, Yung M-H, Yuan X (2021) Low-depth quantum state preparation. *Phys Rev Res* 3(4):043200. <https://doi.org/10.1103/PhysRevResearch.3.043200>
- Zhu C, Byrd RH, Lu P, Nocedal J (1997) Algorithm 778: L-BFGS-B: Fortran subroutines for large-scale bound-constrained optimization. *ACM Trans Math Softw* 23(4):550–560. <https://doi.org/10.1145/279232.279236>

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