Quantum adiabatic machine learning

Kristen L. Pudenz · Daniel A. Lidar

Abstract We develop an approach to machine learning and anomaly detection via quantum adiabatic evolution. This approach consists of two quantum phases, with some amount of classical preprocessing to set up the quantum problems. In the training phase we identify an optimal set of weak classifiers, to form a single strong classifier. In the testing phase we adiabatically evolve one or more strong classifiers on a superposition of inputs in order to find certain anomalous elements in the classification space. Both the training and testing phases are executed via quantum adiabatic evolution. All quantum processing is strictly limited to two-qubit interactions so as to ensure physical feasibility. We apply and illustrate this approach in detail to the problem of software verification and validation, with a specific example of the learning phase applied to a problem of interest in flight control systems. Beyond this example, the algorithm can be used to attack a broad class of anomaly detection problems.

Keywords Adiabatic quantum computation · Quantum algorithms · Software verification and validation · Anomaly detection

K. L. Pudenz (✉)
Department of Electrical Engineering, Center for Quantum Information Science and Technology, University of Southern California, Los Angeles, CA 90089, USA
e-mail: pudenz@usc.edu

D. A. Lidar
Departments of Electrical Engineering, Chemistry, and Physics,
Center for Quantum Information Science and Technology,
University of Southern California, Los Angeles, CA 90089, USA
e-mail: lidar@usc.edu

🕊 Springer

1 Introduction

Machine learning is a field of computational research with broad applications, ranging from image processing to analysis of complex systems such as the stock market. There is abundant literature concerning learning theory in the classical domain, addressing speed and accuracy of the learning process for different classes of concepts [1]. Groundwork for machine learning using quantum computers has also been laid, showing that quantum machine learning, while requiring as much input information as classical machine learning, may be faster and is capable of handling concepts beyond the reach of any classical learner [2,3].

We consider the machine learning problem of binary classification, assigning a data vector to one of two groups based on criteria derived from a set of training examples provided to the algorithm beforehand. The learning method we use is boosting, whereby multiple *weak classifiers* are combined to create a *strong classifier* formula that is more accurate than any of its components alone [4,5]. This method can be applied to any problem where the separation of two groups of data is required, whether it is distinguishing two species of plants based on their measurements or picking out the letter " a" from all other letters of the alphabet when it is scanned. Our approach to classification is based on recent efforts in boosting using adiabatic quantum optimization (AQO) which showed advantages over classical boosting in the sparsity of the classifiers achieved and their accuracy (for certain problems) [6,7].

As a natural outgrowth of the classification problem, we also formulate a scheme for anomaly detection using quantum computation. Anomaly detection has myriad uses, some examples of which are detection of insider trading, finding faults in mechanical systems, and highlighting changes in time-lapsed satellite imagery [8]. Specifically, we pursue the verification and validation (V&V) of classical software, with programming errors as the anomalies to be detected. This is one of the more challenging potential applications of quantum anomaly detection, because programs are large, complex, and highly irregular in their structure. However, it is also an important and currently intractable problem for which even small gains are likely to yield benefits for the software development and testing community.

The complexity of the V&V problem is easily understood by considering the number of operations necessary for an exhaustive test of a piece of software. Covering every possible set of inputs that could be given to the software requires a number of tests that is exponential in the number of input variables, notwithstanding the complexity of each individual test [9]. Although exhaustive testing is infeasible due to its difficulty, the cost of this infeasibility is large—in 2002, NIST estimated that tens of billions of dollars were lost due to inadequate testing [10].

The subject of how to best implement software testing given limited resources has been widely studied. Within this field, efforts focused on combinatorial testing have found considerable success and will be relevant to our new approach. Combinatorial testing focuses on using the test attempts available to test all combinations of up to a small number, t, of variables, with the idea that errors are usually caused by the interaction of only a few parameters [11,12]. This approach has found considerable success [13,14], with scaling that is logarithmic in n, the number of software parameters, and exponential in t.

Currently, the use of formal methods in the coding and verification phases of software development is the only way to guarantee absolute correctness of software without implementing exhaustive testing. However, formal methods are also expensive and time-consuming to implement. Model checking, a method of software analysis which aims to ensure the validity of all reachable program states, solves n-bit satisfiability problems (which are NP-complete), with n as a function of the number of reachable states of the program [15]. Theorem proving, where a program is developed alongside a proof of its own correctness, requires repeated interaction and correction from the developer as the proof is formed, with the intermediate machine-provable lemmas checked with a satisfiability solver [16].

We propose a new approach to verification and validation of software which makes use of quantum information processing. The approach consists of a quantum learning step and a quantum testing step. In the learning step, our strategy uses quantum optimization to learn the characteristics of the program being tested and the specification to which it is being tested. This learning technique is known as quantum boosting and has been previously applied to other problems, in particular image recognition [6,17–19]. Boosting consists of building up a formula to accurately sort inputs into one of two groups by combining simple rules that sort less accurately, and in its classical forms has been frequently addressed in the machine learning literature [4,5,20].

The testing step is novel, and involves turning the classifying formulas generated by the learning step into a function that generates a lower energy when it is more likely that its input represents a software error. This function is translated into the problem Hamiltonian of an adiabatic quantum computation (AQC). The AQC allows all potential software errors (indeed, as we will see, all possible operations of the software) to be examined in quantum-parallel, returning only the best candidates for errors which correspond to the lowest values of the classification function.

Both the learning and testing steps make use of AQC. An adiabatic quantum algorithm encodes the desired result in the ground state of some problem Hamiltonian. The computation is then performed by initializing a physical system in the easily prepared ground state of a simpler Hamiltonian, then slowly changing the control parameters of the system so that the system undergoes an adiabatic evolution to the ground state of the difficult-to-solve problem Hamiltonian [21,22]. The adiabatic model of quantum computation is known to be universal and equivalent to the circuit model with a polynomial conversion overhead [23,24]. While it is not known at this time how to make AQC fault tolerant, several error correction and prevention protocols have been proposed for AQC [25,26], and it is known to exhibit a certain degree of natural robustness [27,28].

In this article, Sect. 2 will begin by establishing the framework through which the quantum V&V problem is attacked, and by defining the programming errors we seek to eliminate. As we proceed with the development of a method for V&V using quantum resources, Sect. 3 will establish an implementation of the learning step as an adiabatic quantum algorithm. We develop conditions for ideal boosting and an alternate quantum learning algorithm in Sect. 4. The testing step will be detailed in Sect. 5. We present simulated results of the learning step on a sample problem in Sect. 6, and finish with our conclusions and suggestions for future work in Sect. 7.

2 Formalization

In this section we formalize the problem of software error detection by first introducing the relevant vector spaces and then giving a criterion for the occurrence of an error.

2.1 Input and output spaces

Consider an " ideal" software program $\hat{P}$, where by ideal we mean the correct program which a perfect programmer would have written. Instead, we are faced with the actual implementation of $\hat{P}$, which we denote by P and refer to as the "implemented program." Suppose we wish to verify the operation of P relative to $\hat{P}$. All programs have input and output spaces $\mathcal{V}_{\text{in}}$ and $\mathcal{V}_{\text{out}}$, such that

$$P : \mathcal{V}_{\text{in}} \mapsto \mathcal{V}_{\text{out}}. \tag{1}$$

Without loss of generality we can think of these spaces as being spaces of binary strings. This is so because the input to any program is always specified within some finite machine precision, and the output is again given within finite machine precision (not necessarily the same as the input precision). Furthermore, since we are only interested in inputs and outputs which take a finite time to generate (or " write down"), without loss of generality we can set upper limits on the lengths of allowed input and output strings. Within these constraints we can move to a binary representation for both input and output spaces, and take N_{in} as the maximum number of bits required to specify any input, and N_{out} as the maximum number of bits required to specify any output. Thus we can identify the input and output spaces as binary string spaces

$$\mathcal{V}_{\text{in}} \cong \{0, 1\}^{N_{\text{in}}}, \quad \mathcal{V}_{\text{out}} \cong \{0, 1\}^{N_{\text{out}}}. \tag{2}$$

It will be convenient to concatenate elements of the input and output spaces into single objects. Thus, consider binary vectors $\mathbf{x} = (\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{out}})$, where $\mathbf{x}_{\text{out}} = P(\mathbf{x}_{\text{in}})$, consisting of program input–output pairs:

$$\mathbf{x} \in \{0, 1\}^{N_{\text{in}}} \times \{0, 1\}^{N_{\text{out}}} = \{0, 1\}^{N_{\text{in}}+N_{\text{out}}} \equiv \mathcal{V}. \tag{3}$$

2.2 Recognizing software errors

2.2.1 Validity domain and range

We shall assume without loss of generality that the input spaces of the ideal and implemented programs are identical. This can always be ensured by extending the ideal program so that it is well defined for all elements of $\mathcal{V}_{\text{in}}$. Thus, while in general not all elements of $\mathcal{V}_{\text{in}}$ have to be allowed inputs into $\hat{P}$ (for example, an input vector that is out of range for the ideal program), one can always reserve some fixed value for such inputs (e.g., the largest vector in $\mathcal{V}_{\text{out}}$) and trivially mark them as errors. The ideal

program $\hat{P}$ is thus a map from the input space to the space $\mathcal{R}_{\text{out}}$ of correct outputs:

$$\hat{P} : \mathcal{V}_{\text{in}} \mapsto \mathcal{R}_{\text{out}} \subseteq \mathcal{V}_{\text{out}}. \tag{4}$$

More specifically, $\hat{P}$ computes an output string $\hat{\vec{x}}_{\text{out}}$ for every input string $\mathbf{x}_{\text{in}}$, i.e., we can write $\hat{\vec{x}}_{\text{out}} = \hat{P}(\mathbf{x}_{\text{in}})$. Of course this map can be many-to-one (non-injective and surjective), but not one-to-many (multi-valued).[1] The implemented program P should ideally compute the exact same function. In reality it may not. With this in mind, the simplest way to identify a software error is to find an input vector $\mathbf{x}_{\text{in}}$ such that

$$\|\hat{P}(\mathbf{x}_{\text{in}}) - P(\mathbf{x}_{\text{in}})\| \neq 0. \tag{5}$$

in some appropriate norm. This is clearly a sufficient condition for an error, since the implemented program must agree with the ideal program on all inputs. However, for our purposes a more general approach will prove to be more suitable.

2.2.2 Specification and implementation sets

A direct way to think about the existence of errors in a software program is to consider two ordered sets within the space of input–output pairs, $\mathcal{V}$. These are the set of ordered, correct input–output pairs $\hat{S}$ according to the program specification $\hat{P}$, and the set of input–output pairs S implemented by the real program P. We call $\hat{S}$ the "specification set" and S the "implementation set". The program under test is correct when

$$\hat{S} = S. \tag{6}$$

That is, in a correct program, the specification set of correct input–output pairs is exactly the set that is implemented in code.

As stated, (6) is impractical since it requires knowledge of the complete structure of the intended input and output spaces. Instead, we can also use the specification and implementation sets to give a correctness criterion for a given input–output pair:

Definition 1 A vector $\mathbf{x} \in \mathcal{V}$ is erroneous and implemented if

$$\mathbf{x} \notin \hat{S} \quad \& \quad \mathbf{x} \in S. \tag{7}$$

Input–output vectors satisfying (7) are the manifestation of software errors ("bugs") and their identification is the main problem we are concerned with here. Conversely, we have

Definition 2 A vector $\mathbf{x} \in \mathcal{V}$ is correct and implemented if

$$\mathbf{x} \in \hat{S} \quad \& \quad \mathbf{x} \in S. \tag{8}$$

---

[1] Random number generation may appear to be a counterexample, as it is multi-valued, but only over different calls to the random-number generator.

4
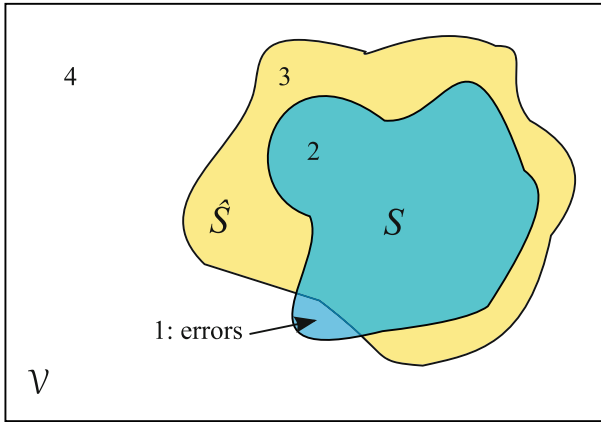
3

2

$\hat{S}$

S

1: errors

$\mathcal{V}$

Fig. 1 Schematic vector space representation showing regions of vectors satisfying the four definitions. *Region 1*, of erroneous but implemented vectors, is the location of errors. *Regions 2, 3*, and *4* represent vectors which are correct and implemented, correct and unimplemented, and erroneous and unimplemented, respectively

Input–output vectors satisfying (8) belong to the "don't-worry" class. The two other possibilities belong to the "don't-care" class:

Definition 3 A vector $\mathbf{x} \in \mathcal{V}$ is correct and unimplemented if

$$\mathbf{x} \in \hat{S} \quad \& \quad \mathbf{x} \notin S. \tag{9}$$

Definition 4 A vector $\mathbf{x} \in \mathcal{V}$ is erroneous and unimplemented if

$$\mathbf{x} \notin \hat{S} \quad \& \quad \mathbf{x} \notin S. \tag{10}$$

A representation of the locations of vectors satisfying the four definitions for a sample vector space can be found in Fig. 1. Our focus will be on the erroneous vectors of Definition 1.

Note that Eq. (5) implies that the vector is erroneous and implemented, i.e., Definition 1. Indeed, let $\mathbf{x}_{\text{in}} = P(\mathbf{x}_{\text{in}})$, i.e., $\mathbf{x} = (\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{out}}) \in S$, but assume that $\mathbf{x}_{\text{out}} \neq \hat{\tilde{x}}_{\text{out}}$ where $\hat{\tilde{x}}_{\text{out}} = \hat{P}(\mathbf{x}_{\text{in}})$. Then $\mathbf{x} \notin \hat{S}$, since $\mathbf{x}_{\text{in}}$ pairs up with $\hat{\tilde{x}}_{\text{out}}$ in $\hat{S}$. Conversely, Definition 1 implies Eq. (5). To see this, assume that $\mathbf{x} = (\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{out}}) \in S$ but $\mathbf{x} = (\mathbf{x}_{\text{in}}, \mathbf{x}_{\text{out}}) \notin \hat{S}$. This must mean that $\mathbf{x}_{\text{out}} \neq \hat{\tilde{x}}_{\text{out}}$, again because $\mathbf{x}_{\text{in}}$ pairs up with $\hat{\tilde{x}}_{\text{out}}$ in $\hat{S}$. Thus Eq. (5) is in fact equivalent to Definition 1, but does not capture the other three possibilities captured by Definitions 2–4.

Definitions 1–4 will play a central role in our approach to quantum V&V.

2.2.3 Generalizations

Note that it may well be advantageous in practice to consider a more general setup, where instead of studying only the map from the input to the output space, we introduce

intermediate maps which track intermediate program states. This can significantly improve our error classification accuracy.[2] Formally, this would mean that Eq. (4) is replaced by

$$\hat{P} : \mathcal{V}_{\text{in}} \mapsto \mathcal{I}_1 \mapsto \cdots \mapsto \mathcal{I}_J \mapsto \mathcal{R}_{\text{out}}, \tag{11}$$

where $\{\mathcal{I}_j\}_{j=1}^J$ are intermediate spaces. However, we shall not consider this more refined approach in this work.

As a final general comment, we reiterate that a solution of the problem we have defined has implications beyond V&V. Namely, Definitions 1–4 capture a broad class of anomaly (or outlier) detection problems [8]. From this perspective the approach we detail in what follows can be described as "quantum anomaly detection", and could be pursued in any application which requires the batch processing of a large data space to find a few anomalous elements.

3 Training a quantum software error classifier

In this section we discuss how to identify whether a given set of input–output pairs is erroneous or correct, and implemented or unimplemented, as per Definitions 1–4. To this end we shall require so-called *weak classifiers*, a *strong classifier*, a methodology to efficiently train the strong classifier, and a way to efficiently apply the trained strong classifier on all possible input–output pairs. Both the training step and the application step will potentially benefit from a quantum speedup.

3.1 Weak classifiers

Consider a class of functions which map from the input–output space to the reals:

$$h_i : \mathcal{V} \mapsto \mathbb{R}. \tag{12}$$

We call these functions "weak classifiers" or "feature detectors," where $i \in \{1, ..., N\}$ enumerates the features. These are some predetermined useful aggregate characteristics of the program P which we can measure, such as total memory, or CPU time average [29]. Note that N will turn out to be the number of qubits we shall require in our quantum approach.

We can now formally associate a weak classification with each vector in the input–output space.

Definition 5 Weak classification of $\mathbf{x} \in \mathcal{V}$.
Weakly classified correct (WCC): a vector $\mathbf{x}$ is WCC if $h_i(\mathbf{x}) > 0$.
Weakly classified erroneous (WCE): a vector $\mathbf{x}$ is WCE if $h_i(\mathbf{x}) < 0$.

---

[2] One important consideration is that, as we shall see below, for practical reasons we may only be able to track errors at the level of one-bit errors and correlations between bit-pairs. Such limited tracking can be alleviated to some extent by using intermediate spaces, where higher order correlations between bits appearing at the level of the output space may not yet have had time to develop.

Clearly, there is an advantage to finding " smart" weak classifiers, so as to minimize N. This can be done by invoking heuristics, or via a systematic approach such as one we present below.

For each input–output pair $\mathbf{x}$ we have a vector $\mathbf{h}(\mathbf{x}) = (h_1(\mathbf{x}), ..., h_N(\mathbf{x})) \in \mathbb{R}^N$. Such vectors can be used to construct geometric representations of the learning problem, e.g., a convex hull encompassing the weak classifier vectors of clustered correct input–output pairs. Such a computational geometry approach was pursued in [29].

We assume that we can construct a " training set"

$$\mathcal{T} \equiv \{\mathbf{x}_s, y_s\}_{s=1}^{|\mathcal{T}|}, \tag{13}$$

where each $\mathbf{x}_s \in \mathcal{V}$ is an input–output pair and $y_s = y(\mathbf{x}_s) = +1$ iff $\mathbf{x}_s$ is correct (whether implemented or not, i.e., $\mathbf{x}_s \in \hat{S}$) while $y_s = -1$ iff $\vec{x}_s$ is erroneous (again, implemented or not, i.e., $\mathbf{x}_s \notin \hat{S}$). Thus, the training set represents the ideal program $\hat{P}$, i.e., we assume that the training set can be completely trusted. Note that Eq. (4) presents us with an easy method for including erroneous input pairs, by deliberately misrepresenting the action of $\hat{P}$ on some given input, e.g., by setting $\mathbf{x}_{\text{out}} \notin \mathcal{R}_{\text{out}}(\hat{P})$. This is similar to the idea of performing V&V by building invariants into a program [30].

We are free to normalize each weak classifier so that $h_i \in [-1/N, 1/N]$ (the reason for this will become clear below). Given Definition 5 we choose the sign of each weak classifier so that $h_i(\mathbf{x}_s) < 0$ for all erroneous training data, while $h_i(\mathbf{x}_s) > 0$ for all correct training data. Each point $\mathbf{h}(\mathbf{x}_s) \in [-1/N, 1/N]^N$ (a hypercube) has associated with it a label y_s which indicates whether the point is correct or erroneous. The convex hull approach to V&V [29] assumes that correct training points $\mathbf{h}(\mathbf{x}_s)$ cluster. Such an assumption is not required in our approach.

3.2 Strong classifier

We would like to combine all the weak classifiers into a single "strong classifier" which, given an input–output pair, will determine that pair's correctness or erroneousness. The problem is that we do not know in advance how to rank the weak classifiers by relative importance. We can formally solve this problem by associating a weight $w_i \in \mathbb{R}$ with each weak classifier h_i. The problem then becomes how to find the optimal set of weights, given the training set.

The process of creating a high-performance strong classifier from many less accurate weak classifiers is known as boosting in the machine learning literature. Boosting is a known method for enhancing to arbitrary levels the performance of known sets of classifiers that exhibit weak learnability for a problem, i.e., they are accurate on more than half of the training set [20,31]. The most efficient method to combine weak classifiers into a strong classifier of a given accuracy is an open question, and there are many competing algorithms available for this purpose [32,33]. Issues commonly considered in the development of such algorithms include identification of the data features that are relevant to the classification problem at hand [34,35] and whether or not provisions need to be taken to avoid overfitting to the training set (causing poor

performance on the general problem space) [36,37]. We use an approach inspired by recent quantum boosting results on image recognition [6,17–19]. This approach has been shown to outperform classical boosting algorithms in terms of accuracy (but not speed) on selected problems, and has the advantage of being implementable on existing quantum optimization hardware [38–41].

Since we shall map the w_i to qubits we use binary weights $w_i \in \{0, 1\}$. It should be straightforward to generalize our approach to a higher resolution version of real-valued w_i using multiple qubits per weight.

Let $\mathbf{w} = (w_1, ..., w_N) \in \{0, 1\}^N$, and let

$$R_{\mathbf{w}}(\mathbf{x}) \equiv \mathbf{w} \cdot \mathbf{h}(\mathbf{x}) = \sum_{i=1}^{N} w_i h_i(\mathbf{x}) \in [-1, 1]. \qquad (14)$$

This range is a direct result of the normalization $h_i \in [-1/N, 1/N]$ introduced above.

We now define the weight-dependent " strong classifier"

$$Q_{\mathbf{w}}(\mathbf{x}) \equiv \text{sign}\,[R_{\mathbf{w}}(\mathbf{x})], \qquad (15)$$

and use it as follows:

Definition 6 Strong classification of $\mathbf{x} \in \mathcal{V}$.
Strongly classified correct (SCC): a vector $\mathbf{x}$ is SCC if $Q_{\mathbf{w}}(\mathbf{x}) = +1$.
Strongly classified erroneous (SCE): a vector $\mathbf{x}$ is SCE if $Q_{\mathbf{w}}(\mathbf{x}) = -1$.

There is a fundamental difference between the " opinions" of the strong classifier, as expressed in Definition 6, and the actual erroneousness/correctness of a given input–output pair. The strong classifier associates an erroneous/correct label with a given input–output pair according to a weighted average of the weak classifiers. This opinion may or may not be correct. For the training set we actually know whether a given input–output pair is erroneous or correct. This presents us with an opportunity to compare the strong classifier to the training data. Namely, if $y_s Q_{\mathbf{w}}(\vec{x}_s) = -1$ then $Q_{\mathbf{w}}(\mathbf{x}_s)$ and y_s have opposite sign, i.e., disagree, which means that $Q_{\mathbf{w}}(\mathbf{x}_s)$ mistakenly classified $\mathbf{x}_s$ as a correct input–output pair while in fact it was erroneous, or vice versa. On the other hand, if $y_s Q_{\mathbf{w}}(\mathbf{x}_s) = +1$ then $Q_{\mathbf{w}}(\mathbf{x}_s)$ and y_s agree, which means that $Q_{\mathbf{w}}(\mathbf{x}_s)$ is correct. Formally,

$$y_s Q_{\mathbf{w}}(\mathbf{x}_s) = +1 \iff \begin{cases} (\mathbf{x}_s \text{ is SCC}) = \text{true or} \\ (\mathbf{x}_s \text{ is SCE}) = \text{true} \end{cases} \qquad (16a)$$

$$y_s Q_{\mathbf{w}}(\mathbf{x}_s) = -1 \iff \begin{cases} (\mathbf{x}_s \text{ is SCC}) = \text{false or} \\ (\mathbf{x}_s \text{ is SCE}) = \text{false} \end{cases} \qquad (16b)$$

The higher the number of true instances is relative to the number of false instances, the better the strong classifier performance over the training set. The challenge is, of course, to construct a strong classifier that performs well also beyond the training set. To do so we must first solve the problem of finding the optimal set of binary weights $\mathbf{w}$.

3.3 The formal weight optimization problem

Let $\mathbf{H}[z]$ denote the Heaviside step function, i.e., $\mathbf{H}[z] = 0$ if $z < 0$ and $\mathbf{H}[z] = 1$ if $z \geq 0$. Thus $\mathbf{H}[-y_s Q_{\mathbf{w}}(\mathbf{x}_s)] = 1$ if the classification of $\mathbf{x}_s$ is wrong, but $\mathbf{H}[-y_s Q_{\mathbf{w}}(\mathbf{x}_s)] = 0$ if the classification of $\vec{x}_s$ is correct. In this manner $\mathbf{H}[-y_s Q_{\mathbf{w}}(\mathbf{x}_s)]$ assigns a penalty of one unit for each incorrectly classified input–output pair.

Consider

$$L(\mathbf{w}) \equiv \sum_{s=1}^{|\mathcal{T}|} \mathbf{H}[-y_s Q_{\mathbf{w}}(x_s)]. \tag{17}$$

This counts the total number of incorrect classifications. Therefore minimization of $L(\mathbf{w})$ for a given training set $\{\mathbf{x}_s, y_s\}_{s=1}^{|\mathcal{T}|}$ will yield the optimal set of weights $\mathbf{w}^{\text{opt}} = \{w_i^{\text{opt}}\}_{i=1}^{N}$.

However, it is important not to overtrain the classifier. Overtraining means that the strong classifier has poor generalization performance, i.e., it does not classify accurately outside of the training set [37,42]. To prevent overtraining we can add a penalty proportional to the Hamming weight of $\mathbf{w}$, i.e., to the number of non-zero weights $\|\mathbf{w}\|_0 = \sum_{i=1}^{N} w_i$. In this manner an optimal balance is sought between the accuracy of the strong classifier and the number of weak classifiers comprising the strong classifier. The formal weight optimization problem is then to solve

$$\mathbf{w}'^{opt} = \arg \min_{\mathbf{w}} [L(w) + \lambda \|\mathbf{w}\|_0], \tag{18}$$

where $\lambda > 0$ can be tuned to decide the relative importance of the penalty.

3.4 Relaxed weight optimization problem

Unfortunately, the formulation of (18) is unsuitable for adiabatic quantum computation because of its discrete nature. In particular, the evaluation of the Heaviside function is not amenable to a straightforward implementation in AQC. Therefore, following [6], we now relax it by introducing a quadratic error measure, which will be implementable in AQC.

Let $\mathbf{y} = (y_1, ..., y_{|\mathcal{T}|}) \in \{-1, 1\}^{|\mathcal{T}|}$ and $\mathbf{R}_{\mathbf{w}} = (R_{\mathbf{w}}(\mathbf{x}_1), ..., R_{\mathbf{w}}(\mathbf{x}_{|\mathcal{T}|})) \in [-1, 1]^{|\mathcal{T}|}$. The vector $\mathbf{y}$ is the ordered label set of correct/erroneous input–output pairs. The components $R_{\mathbf{w}}(\mathbf{x})$ of the vector $\mathbf{R}_{\mathbf{w}}$ already appeared in the strong classifier (15). There we were interested only in their signs and in Eq. (16) we observed that if $y_s R_{\mathbf{w}}(\mathbf{x}_s) < 0$ then $\mathbf{x}_s$ was incorrectly classified, while if $y_s R_{\mathbf{w}}(\mathbf{x}_s) > 0$ then $\mathbf{x}_s$ was correctly classified.

We can consider a relaxation of the formal optimization problem (18) by replacing the counting of incorrect classifications by a sum of the values of $y_s R_{\mathbf{w}}(\mathbf{x}_s)$ over the training set. This seems reasonable since we have normalized the weak classifiers so that $R_{\mathbf{w}}(\mathbf{x}) \in [-1, 1]$, while each label $y_s \in \{-1, 1\}$, so that all the

terms $y_s R_{\mathbf{w}}(\mathbf{x}_s)$ are in principle equally important. In other words, the inner product $\mathbf{y} \cdot \mathbf{R_w} = \sum_{s=1}^{|\mathcal{T}|} y_s R_{\mathbf{w}}(\mathbf{x}_s)$ is also a measure of the success of the classification, and maximizing it (making $\mathbf{y}$ and $\mathbf{R_w}$ as parallel as possible) should result in a good training set.

Equivalently, we can consider the distance between the vectors $\mathbf{y}$ and $\mathbf{R_w}$ and minimize it by finding the optimal weight vector $\mathbf{w}^{\text{opt}}$, in general different from that in Eq. (18). Namely, consider the Euclidean distance

$$
\delta(\mathbf{w}) = \|\mathbf{y} - \mathbf{R_w}\|^2 = \sum_{s=1}^{|\mathcal{T}|} \left| y_s - \sum_{i=1}^{N} w_i h_i(x_s) \right|^2
$$

$$
= \|\mathbf{y}\|^2 + \sum_{i,j=1}^{N} C'_{ij} w_i w_j - 2 \sum_{i=1}^{N} C'_{iy} w_i, \tag{19}
$$

where $\mathbf{h}_i = (h_i(x_1), ..., h_i(x_{|\mathcal{T}|})) \in [-1/N, 1/N]^{|\mathcal{T}|}$ and where

$$
C'_{ij} = \mathbf{h}_i \cdot \mathbf{h}_j = \sum_{s=1}^{|\mathcal{T}|} h_i(x_s) h_j(x_s), \tag{20}
$$

$$
C'_{iy} = \mathbf{h}_i \cdot \mathbf{y} = \sum_{s=1}^{|\mathcal{T}|} h_i(x_s) y_s \tag{21}
$$

can be thought of as correlation functions. Note that they are symmetric: $C'_{ij} = C'_{ji}$ and $C'_{iy} = C'_{yi}$. The term $\|\mathbf{y}\|^2 = |\mathcal{T}|$ is a constant offset and can be dropped from the minimization.

If we wish to introduce a sparsity penalty as above, we can do so again, and thus ask for the optimal weight in the following sense:

$$
\mathbf{w}^{\text{opt}} = \arg\min_{\mathbf{w}} \left[\delta(\mathbf{w}) + \lambda' \|\mathbf{w}\|_0 \right]
$$

$$
= \arg\min_{\mathbf{w}} \left[\sum_{i,j=1}^{N} C'_{ij} w_i w_j + 2 \sum_{i=1}^{N} (\lambda - C'_{iy}) w_i \right], \tag{22}
$$

where $\lambda' = 2\lambda$.

3.5 From QUBO to the Ising Hamiltonian

Equation (22) is a quadratic binary optimization (QUBO) problem [17]. One more step is needed before we can map it to qubits, since we need to work with optimization variables whose range is $\{-1, 1\}$, not $\{0, 1\}$. Define new variables $q_i = 2(w_i - 1/2) \in \{-1, 1\}$. In terms of these new variables the minimization problem is

$$\mathbf{q}^{\text{opt}} = \arg\min_{\mathbf{q}} \left[\frac{1}{4} \sum_{i,j=1}^{N} C'_{ij}(q_i + 1)(q_j + 1) + \sum_{i=1}^{N} (\lambda - C'_{iy})(q_i + 1) \right]$$

$$= \arg\min_{\mathbf{q}} \left[\sum_{i,j=1}^{N} C_{ij} q_i q_j + \sum_{i=1}^{N} (\lambda - C_{iy}) q_i \right], \tag{23}$$

where in the second line we dropped the constant terms $\frac{1}{4} \sum_{i,j=1}^{N} C'_{ij}$ and $\sum_{i=1}^{N} (\lambda - C'_{iy})$, used the symmetry of C'_{ij} for $\sum_{i=1}^{N} q_i \sum_{j=1}^{N} C'_{ij} = \sum_{i,j=1}^{N} C'_{ij} q_j$, and where we defined

$$C_{ij} = \frac{1}{4} C'_{ij}, \quad C_{iy} = C'_{iy} - \frac{1}{2} \sum_{j=1}^{N} C'_{ij}. \tag{24}$$

Thus, the final AQC Hamiltonian for the quantum weight-learning problem is

$$H_F = \sum_{i,j=1}^{N} C_{ij} Z_i Z_j + \sum_{i=1}^{N} (\lambda - C_{iy}) Z_i, \tag{25}$$

where Z_i is the Pauli spin-matrix σ_z acting on the ith qubit. This represents Ising spin-spin interactions with coupling matrix C_{ij}, and an inhomogeneous magnetic field $\lambda - C_{iy}$ acting on each spin. Note how H_F encodes the training data $\{h_i(x_s), y_s\}_{i,s}$ via the coupling matrix $C_{ij} = \frac{1}{4} \sum_{s=1}^{|T|} h_i(x_s) h_j(x_s)$ and the local magnetic field $C_{iy} = \sum_{s=1}^{|T|} h_i(x_s) y_s - \frac{1}{2} \sum_{s=1}^{|T|} h_i(x_s) \sum_{j=1}^{N} h_j(x_s)$. Thus, in order to generate H_F one must first calculate the training data using the chosen set of weak classifiers.

In this final form (Eq. 25), involving only one and two-qubit Z_i terms, the problem is now suitable for implementation on devices such as D-Wave's adiabatic quantum optimization processor [19,39].

In Sect. 4.4 we shall formulate an alternative weight optimization problem, based on a methodology we develop in Sect. 4 for pairing weak classifiers to guarantee the correctness of the strong classifier.

3.6 Adiabatic quantum computation

The adiabatic quantum algorithm implements the time-dependent interpolation

$$H(t) = s(t) H_I + [1 - s(t)] H_F, \tag{26}$$

where H_I is a Hamiltonian which does not commute with H_F and should have a ground state (lowest-energy eigenvector) that is easily reachable, such as

$$H_I = \mathbb{I} - \sum_{i=1}^{N} X_i \tag{27}$$

where $\mathbb{I}$ is the identity operator and X_i is the Pauli σ_x acting on the ith qubit [21,22]. The interpolation function $s(t)$ satisfies the boundary conditions $s(0) = 1$, $s(T) = 0$, where T is the final time. Provided the evolution is sufficiently slow (in a manner we shall quantify momentarily), the adiabatic theorem guarantees that the final state $|\psi(T)\rangle$ reached by the algorithm is, with high probability, the one that minimizes the energy of H_F [43–45]. This means that, for H_F chosen as in Eq. (25), it finds as a ground state the optimal weights vector $\mathbf{q}^{\mathrm{opt}}$ as defined in (23). These weights can then be "read off" by measuring the final states of each of the N qubits: $|\psi(T)\rangle = |q_1^{\mathrm{opt}}, ..., q_N^{\mathrm{opt}}\rangle = |\mathbf{q}^{\mathrm{opt}}\rangle$.

It should be noted that while the number of weak classifiers that can be selected from using this algorithm may appear to be limited by the number of qubits available for processing, this is not in fact the case. By performing multiple rounds of optimization, each time filling in the spaces left by classifiers that were assigned weight 0 in the previous round, an optimized group of N weak classifiers can be assembled. If the performance of the strong classifier is unsatisfactory with N weak classifiers, multiple groups of N found in this manner may be used together.

The scaling of the computation time t_F with the number of qubits (or weak classifiers, in our case), N, is determined by the inverse of the minimal ground state energy gap of $H(t)$. There are many variants of the adiabatic theorem, differing mostly in assumptions about boundary conditions and differentiability of $H(t)$. Most variants state that, provided

$$t_F \gtrsim \frac{\|\dot{H}\|^\alpha}{\epsilon \Delta^{\alpha+1}}, \tag{28}$$

then

$$|\langle \psi(t_F)|\phi(t_F)\rangle| \gtrsim 1 - \epsilon^\beta. \tag{29}$$

The left-hand side of Eq. (29) is the fidelity of the *actual* state $|\psi(t_F)\rangle$ obtained under quantum evolution subject to $H(t)$ with respect to the *desired* final ground state $|\phi(t_F)\rangle$. More precisely, $|\psi(t)\rangle$ is the solution of the time-dependent Schrödinger equation $\partial|\psi(t)\rangle/\partial t = -i H(t)|\psi(t)\rangle$ (in $\hbar \equiv 1$ units), and $|\phi(t)\rangle$ is the instantaneous ground state of $H(t)$, i.e., the solution of $H(t)|\phi(t)\rangle = E_0(t)|\phi(t)\rangle$, where $E_0(t)$ is the instantaneous ground state energy [the smallest eigenvalue of $H(t)$]. The parameter ϵ, $0 \le \epsilon \le 1$, measures the quality of the overlap between $|\psi(t_F)\rangle$ and $|\phi(t_F)\rangle$, $\dot{H}$ is the derivative with respect to the dimensionless time t/t_F, Δ is the minimum energy gap between the ground state $|\phi(t)\rangle$ and the first excited state of $H(t)$ (i.e., the difference between the two smallest equal-time eigenvalues of $H(t)$, for $t \in [0, t_F]$), The values of the integers α and β depend on the assumptions made about the boundary conditions and differentiability of $H(t)$ [43–45]; typically $\alpha \in \{0, 1, 2\}$, while β can be tuned between 1 and arbitrarily large values, depending on boundary conditions determining the smoothness of $H(t)$ (see, e.g., Theorem 1 in Ref. [45]). The crucial point is that the gap Δ depends on N, typically shrinking as N grows, while the numerator $\|\dot{H}\|$ typically has a mild N-dependence (bounded in most cases by a function growing as N^2 [45]). Consequently a problem has an efficient, polynomial time solution under

AQC if Δ scales $1/\text{poly}(N)$. However, note that an inverse exponential gap dependence on N can still result in a speedup, as is the case, e.g., in the adiabatic implementation of Grover's search problem [46,47], where the speedup relative to classical computation is quadratic.

As for the problem we are concerned with here, finding the ground state of H_F as prescribed in Eq. (25) in order to find the optimal weight set for the (relaxed version of the) problem of training a software error-classifier, it is not known whether it is amenable to a quantum speedup. A study of the gap dependence of our Hamiltonian $H(t)$ on N, which is beyond the scope of the present work, will help to determine whether such a speedup is to be expected also in the problem at hand. A related image processing problem has been shown numerically to require fewer weak classifiers than in comparable classical algorithms, which gives the strong classifier a lower Vapnik-Chernovenkis dimension and therefore a lower generalization error [7,18]. Quantum boosting applied to a different task, 30-dimensional clustering, demonstrated increasingly better accuracy as the overlap between the two clusters grew than that exhibited by the classical AdaBoost algorithm [6]. More generally, numerical simulations of quantum adiabatic implementations of related hard optimization problems (such as Exact Cover) have shown promising scaling results for N values of up to 128 [22,48,40]. We shall thus proceed here with the requisite cautious optimism.

4 Achievable strong classifier accuracy

We shall show in this section that it is theoretically possible to construct a perfect, 100 % accurate majority-vote strong classifier from a set of weak classifiers that are more than 50 % accurate—if those weak classifiers relate to each other in exactly the right way. Our construction in this section is analytical and exact; we shall specify a set of conditions weak classifiers should satisfy for perfect accuracy of the strong classifier they comprise. We shall also show how to construct an imperfect strong classifier, with bounded error probability, by a relaxation of the conditions we shall impose on the weak classifiers. We expect the quantum algorithm to find a close approximation to this result.

Consider a strong classifier with a general binary weight vector $\mathbf{w} \in \{0, 1\}^N$, as defined in Eq. (14). Our approach will be to show that the strong classifier in Eq. (14) is completely accurate if a set of three conditions is met. The conditions work by using pairs of weak classifiers which both classify some $\mathbf{x}$ correctly and which disagree for all other $\mathbf{x}$. An accurate strong classifier can be constructed by covering the entire space $\mathcal{V}$ with the correctly classifying portions of such weak classifier pairs.

To start, every vector $\mathbf{x} \in \mathcal{V}$ has a correct classification, as determined by the specification set:

$$\mathbf{x} \in \hat{S} \iff y(\mathbf{x}) = +1, \tag{30a}$$
$$\mathbf{x} \notin \hat{S} \iff y(\mathbf{x}) = -1 \tag{30b}$$

A strong classifier is perfect if

$$Q_{\mathbf{w}}(\mathbf{x}) = y(\mathbf{x}) \quad \forall \mathbf{x} \in \mathcal{V}. \tag{31}$$

The weak classifiers either agree or disagree with this correct classification. We define the correctness value of a weak classifier for a given input $\mathbf{x}$:

$$c_i(\mathbf{x}) = h_i(\mathbf{x})y(\mathbf{x}) = \begin{cases} +1 & h_i(\mathbf{x}) = y(\mathbf{x}) \\ -1 & h_i(\mathbf{x}) \neq y(\mathbf{x}) \end{cases} \tag{32}$$

Thus, similarly to the strong classifier case (Eq. (16)) we have, formally,

$$c_i(\mathbf{x}) = +1 \iff \begin{cases} (\mathbf{x} \text{ is WCC}) = \text{true or} \\ (\mathbf{x} \text{ is WCE}) = \text{true} \end{cases} \tag{33a}$$

$$c_i(\mathbf{x}) = -1 \iff \begin{cases} (\mathbf{x} \text{ is WCC}) = \text{false or} \\ (\mathbf{x} \text{ is WCE}) = \text{false} \end{cases} \tag{33b}$$

where WCC and WCE stand for weakly classified correct and weakly classified erroneous, respectively (Definition 5).

A given input–output vector $\mathbf{x}$ receives either a true or false vote from each weak classifier comprising the strong classifier. Let us denote the index set of the weak classifiers comprising a given strong classifier by $\mathcal{I}$. If the majority of the votes given by the weak classifiers in $\mathcal{I}$ are true then the vector receives a strong classification that is true. Let us loosely denote by $\mathbf{w} \in \mathcal{I}$ the set of weak classifiers whose indices all belong to $\mathcal{I}$. Thus

$$\sum_{i \in \mathcal{I}} c_i(\mathbf{x}) > 0 \implies Q_{\mathbf{w}}(\mathbf{x}) = y(\mathbf{x}) \quad \text{if } \mathbf{w} \in \mathcal{I}. \tag{34}$$

It follows from Eq. (31) that if we can find a set of weak classifiers for which $\sum_{i \in \mathcal{I}} c_i(\mathbf{x}) > 0$ for all input–output vectors $\mathbf{x}$, then the corresponding strong classifier is perfect. This is what we shall set out to do in the next subsection.

4.1 Conditions for complete classification accuracy

First, we limit our working set to those weak classifiers with greater than 50 % accuracy. This is a prerequisite for the feasibility of the other conditions. To ensure that at least half the initial dictionary of weak classifiers is more than 50 % accurate, we include each potential weak classifier in the dictionary, as well as its opposite. The opposite classifier follows the same rule as its counterpart, but makes the opposite binary decision every time, making each right where the other is wrong and ensuring that at least one of them will have 50 % or greater accuracy. Condition 1, therefore, defines the set $\mathcal{A}$,

$$\mathcal{A} \subseteq \mathcal{D} \equiv \{1, ..., N\}, \tag{35}$$

of sufficiently accurate weak classifiers, where $\mathcal{D}$ is the set of all possible values of the index i of weak classifiers in Eq. (14).

Condition 1 *For an input–output vector* $\mathbf{x} \in \mathcal{V}$ *selected uniformly at random*

$$\mathcal{A} = \{i : P[c_i(\mathbf{x}) = 1] > 1/2\}. \tag{36}$$

$P[\omega]$ denotes the probability of event ω. We use a probabilistic formulation for our conditions since we imagine the input–output space $\mathcal{V}$ to be very large and accessed by random sampling.

Conditions 2 and 3 (or 3a) specify the index set

$$\mathcal{J} \subseteq \mathcal{A} \times \mathcal{A}, \tag{37}$$

labeling pairs of weak classifiers which will make up the final strong classifier. Condition 2 groups the weak classifiers into pairs which classify the minimal number of vectors $\mathbf{x}$ correctly at the same time and give opposite classifications on all other vectors. Condition 3 completes the specification of the index set $\mathcal{J}$: it states that the subsets of vectors $\mathbf{x}$ that are classified correctly by the classifier pairs in $\mathcal{J}$ must cover the entire space $\mathcal{V}$.

Condition 2 *If* $(j, j') \in \mathcal{J}$ *then*

$$P\left[(c_j(\mathbf{x}) = 1) \cap (c_{j'}(\mathbf{x}) = 1)\right] = P\left[c_j(\mathbf{x}) = 1\right] + P\left[c_{j'}(\mathbf{x}) = 1\right] - 1 \tag{38}$$

for an input–output vector $\mathbf{x} \in \mathcal{V}$ *selected uniformly at random.*

This condition has the following simple interpretation, illustrated in Fig. 2. Suppose the entire input–output space $\mathcal{V}$ is sorted lexicographically (e.g., according to the binary values of the vectors $\mathbf{x} \in \mathcal{V}$) so that the jth weak classifier is correct on all first N_j vectors but erroneous on the rest, while the j'th weak classifier is correct on all last $N_{j'}$ vectors but erroneous on the rest. Thus the fraction of correctly classified vectors by the jth classifier is $(1 - \eta_j) = N_j/|\mathcal{V}|$, the fraction of correctly classified vectors by the j'th classifier is $(1 - \eta_{j'}) = N_{j'}/|\mathcal{V}|$, and they overlap on a fraction of $1 - \eta_j - \eta_{j'}$ vectors (all vectors minus each classifier's fraction of incorrectly classified vectors), as illustrated in the top part of Fig. 2. By "pushing classifier j' to the left", as illustrated in the bottom part of Fig. 2, the overlap grows and is no longer minimal. This is what is expressed by Eq. (38).

Condition 2 considers only one pair of weak classifiers at a time, which does not suffice to cover all of $\mathcal{V}$. Consider a set of weak classifier pairs each satisfying Condition 2 which, together, do cover all of $\mathcal{V}$. Such a set would satisfy

$$\sum_{(j,j')\in\mathcal{J}} P\left[(c_j(\mathbf{x}) = 1) \cap (c_{j'}(\mathbf{x}) = 1)\right] = 1$$

for a randomly chosen $\mathbf{x} \in \mathcal{V}$. This is illustrated in Fig. 3. However, it is also possible for two or more pairs to overlap, a situation we would like to avoid as much as possible, i.e., we shall impose minimal overlap similarly to Condition 2. Thus we arrive at:

Fig. 2 Illustration of Condition 2. Two pairs of classifiers showing regions of correct (*green*) and incorrect (*red*) classification along a *line* representing a lexicographical ordering of all vectors within $\mathcal{V}$. The *top* pair, compliant with Condition 2, provides two correct classifications for the minimum possible number of vectors, voting once correctly and once incorrectly on all other vectors. The *bottom* pair, violating Condition 2, provides two correct votes for more vectors than does the *top* pair, but also undesirably provides two incorrect votes for some vectors; this is why paired weak classifiers must coincide in their classifications on as few vectors as possible (Color figure online)


Fig. 3 Illustration of Condition 3 without the subtracted term. Five pairs of 60 % accurate weak classifiers combine to form a completely accurate majority-vote strong classifier. Moving from *top* to *bottom* through the pairs and from *left* to *right* along the vectors in the classification space, each pair of weak classifiers provides two correct votes for 20 % of the vector space and neutral votes otherwise. This means that the majority vote is correct for the entire space because no two pairs vote correctly at once

Condition 3

$$\sum_{(j,j')\in\mathcal{J}} P\left[(c_j(\mathbf{x})=1)\cap(c_{j'}(\mathbf{x})=1)\right] - \sum_{(j,j')\neq(k,k')\in\mathcal{J}} P\left[(c_j(\mathbf{x})=1)\right.$$
$$\left.\cap(c_{j'}(\mathbf{x})=1)\cap(c_k(\mathbf{x})=1)\cap(c_{k'}(\mathbf{x})=1)\right]=1, \tag{39}$$

where the overlap between two pairs of weak classifiers with labels (j, j') and (k, k') is given by the subtracted terms. Condition 3 is illustrated in Fig. 4.

It is possible to substitute a similar Condition 3a for the above Condition 3 to create a different, yet also sufficient set of conditions for a completely accurate strong classifier. The number of weak classifiers required to satisfy the alternate set of conditions is expected to be smaller than the number required to satisfy the original three conditions.

Fig. 4 Illustration of Condition 3 with the subtracted term. Three pairs of 70 % accurate weak classifiers combined to form a completely accurate majority-vote strong classifier. In this case, each pair votes twice correctly on 40 % of the vector space, which makes it necessary for the correct portions of the second and third pairs from the *top* to overlap. Because they only overlap by the minimum amount necessary, $\mathcal{V}$ as a whole is still covered by a correct majority vote

This is due to the fact that the modified conditions make use of one standalone weak classifier to cover a larger portion of the space correctly than is possible with a pair of weak classifiers.

Condition 3a

$$
\sum_{(j,j')\in\mathcal{J}} P\left[\left(c_j(\mathbf{x}) = 1\right) \cap \left(c_{j'}(\mathbf{x}) = 1\right)\right] + P\left[c_a(\mathbf{x}) = 1\right]
$$
$$
- \sum_{(j,j')\neq(k,k')\in\mathcal{J}} P\left[\left(c_j(\mathbf{x}) = 1\right) \cap \left(c_{j'}(\mathbf{x}) = 1\right) \cap \left(c_k(\mathbf{x}) = 1\right) \cap \left(c_{k'}(\mathbf{x}) = 1\right)\right]
$$
$$
- \sum_{(j,j')\in\mathcal{J}} P\left[\left(c_a(\mathbf{x}) = 1\right) \cap \left(c_j(\mathbf{x}) = 1\right) \cap \left(c_{j'}(\mathbf{x}) = 1\right)\right] = 1 \tag{40}
$$

This condition is illustrated in Fig. 5. Its interpretation is similar to that of Condition 3, except that the standalone classifier with the subscript a is added to the other classifier pairs, and its overlap with them is subtracted separately in the last line.

The perfect strong classifier can now be constructed from the weak classifiers in the set $\mathcal{J}$ defined by the conditions above. Define $\mathcal{J}_L$ as the set of all j from pairs $(j, j') \in \mathcal{J}$. Similarly, define $\mathcal{J}_R$ as the set of all j' from pairs $(j, j') \in \mathcal{J}$. Note that, since any pair for which $j = j'$ would not have minimum correctness overlap and therefore could not be in $\mathcal{J}$, it follows that $j \neq j'$ for all pairs (j, j'), i.e., $\mathcal{J}_L \cap \mathcal{J}_R = \emptyset$. The strong classifier is then (14) with each w_i being one of the elements of a pair, i.e.,

$$
w_i = \begin{cases} 1 & i \in (\mathcal{J}_L \cup \mathcal{J}_R) \\ 0 & \text{otherwise} \end{cases} \tag{41}
$$

4.2 Perfect strong classifier theorem

We will now prove that any strong classifier satisfying Conditions 1–3, or 1–3a, is completely accurate.

Fig. 5 Illustration of Condition 3a. Two pairs and one single weak classifier form a completely accurate majority-vote strong classifier. The two pairs cover 40% of the vector space with correct votes, and the single weak classifier (the first element of the fourth pair in Fig. 3; the faded-out classifiers in the third, fourth, and fifth pairs are omitted from this strong classifier) provides an extra correct vote to tip the balance in the remaining 60% to a correct overall classification

Lemma 1 *Assume Condition 1 and $(j, j') \in \mathcal{J}$. Then the sum of the correctness values of the corresponding weak classifiers is nonnegative everywhere with probability 1, namely*

$$P\left[c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \geq 0\right] = 1 \tag{42}$$

for an input–output vector $\mathbf{x} \in \mathcal{V}$ selected uniformly at random.

Proof For any pair $(j, j') \in \mathcal{J}$ we have

$$P\left[\left(c_j(\mathbf{x}) = 1\right) \cup \left(c_{j'}(\mathbf{x}) = 1\right)\right] = P\left[c_j(\mathbf{x}) = 1\right] + P\left[c_{j'}(\mathbf{x}) = 1\right]$$
$$- P\left[\left(c_j(\mathbf{x}) = 1\right) \cap \left(c_{j'}(\mathbf{x}) = 1\right)\right] = 1 \tag{43}$$

by Condition 2. Equation (43) means that at least one of the two weak classifiers evaluates to 1. Since by definition $c_i(\mathbf{x}) \in \{-1, 1\} \forall i$, the sum is 2 or 0 with probability 1, i.e.,

$$P\left[c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \in \{0, 2\}\right] = 1. \tag{44}$$

Recall that if the majority of the votes given by the weak classifiers comprising a given strong classifier is true then the input–output vector being voted on receives a strong classification that is true (Eq. 34), and that if this is the case for all input–output vectors then the strong classifier is perfect (Eq. 31). We are now in a position to state that this is the case with certainty provided the weak classifiers belong to the set $\mathcal{J}$ defined by the conditions given above.

Theorem 1 *A strong classifier comprised solely of a set of weak classifiers satisfying Conditions 1–3 is perfect.*

Proof It suffices to show that the correctness sum is at least 2 with probability 1 when Conditions 1–3 are met, namely that

$$P\left[\sum_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})\right)\geq 2\right]=1. \tag{45}$$

Now,

$$P\left[\bigcup_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})=2\right)\right]$$

$$=P\left[\bigcup_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right] \tag{46a}$$

$$\geq\sum_{(j,j')\in\mathcal{J}}P\left[\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right]$$

$$-\sum_{(j,j')\neq(k,k')\in\mathcal{J}}P\left[\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right.$$

$$\left.\cap\left(c_k(\mathbf{x})=1\right)\cap\left(c_{k'}(\mathbf{x})=1\right)\right] \tag{46b}$$

$$=1\quad\text{by Cond. 3.} \tag{46c}$$

where equality (46c) holds for the inequality (46b)[3] because the probability of an event cannot be greater than 1.

Thus, for any randomly selected vector $\mathbf{x}\in\mathcal{V}$, the correctness sum of at least one of the pairs is 2, i.e.,

$$P\left[\exists(j,j')\in\mathcal{J}:\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})=2\right)\right]=1. \tag{47}$$

Lemma 1 tells us that the correctness sum of each pair of weak classifiers is positive, while Eq. (47) states that for at least one pair this sum is not just positive but equal to 2. Therefore the correctness sum of all weak classifiers in $\mathcal{J}$ is at least 2, which is Eq. (45).

Theorem 2 *A strong classifier comprised solely of a set of weak classifiers satisfying Conditions 1, 2, and 3a is perfect.*

---

[3] This inequality reflects the fact that for n overlapping sets, $P\left[\bigcup_{i=1}^n s_i\right]=\sum_{i=1}^n P[s_i]-\sum_{i\neq j}P[s_i\cap s_j]+\sum_{i\neq j\neq k}P[s_i\cap s_j\cap s_k]-\sum_{i\neq j\neq k\neq m}P[s_i\cap s_j\cap s_k\cap s_m]+\dots$ Each term is larger than the next in the series; $n+1$ sets cannot intersect where n sets do not. Our truncation of the series is greater than or equal to the full value because we stop after a subtracted term.

Proof It suffices to show that the correctness sum is at least 1 with probability 1 when Conditions 1, 2, and 3a are met, namely that

$$P\left[\sum_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})\right)+c_a(\mathbf{x})\geq 1\right]=1. \tag{48}$$

We proceed similarly to the proof of Theorem 1.

$$P\left[\bigcup_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})=2\right)\cup(c_a(\mathbf{x})=1)\right]$$

$$=P\left[\bigcup_{(j,j')\in\mathcal{J}}\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\cup(c_a(\mathbf{x})=1)\right]$$

$$=\sum_{(j,j')\in J}P\left[\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right]+P\left[c_a(\mathbf{x})=1\right]$$

$$-\sum_{(j,j')\neq(k,k')\in\mathcal{J}}P\left[\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\cap(c_k(\mathbf{x})=1)\cap(c_{k'}(\mathbf{x})=1)\right]$$

$$-\sum_{(j,j')\in\mathcal{J}}P\left[(c_a(\mathbf{x})=1)\cap\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right]$$

$$=1\quad\text{by Cond. 3.} \tag{49}$$

Thus the correctness sum of at least one of the pairs together with the singled-out weak classifier is greater than or equal to 1, i.e.,

$$P\left[\exists(j,j')\in\mathcal{J}:\left(c_j(\mathbf{x})+c_{j'}(\mathbf{x})=2\right)\cup(c_a(\mathbf{x})=1)\right]=1. \tag{50}$$

This result, together with Lemma 1, implies the correctness sum of all weak classifiers in $\mathcal{J}$ is at least 1, which is Eq. (48).

4.3 Imperfect strong classifier theorem

Because the three conditions on the set $\mathcal{J}$ of weak classifiers guarantee a completely accurate strong classifier, errors in the strong classifier must mean that the conditions are violated in some way. For instance, Condition 2 could be replaced by a weaker condition which allows for more than minimum overlap of vectors $\mathbf{x}$ categorized correctly by both weak classifiers in a pair.

Condition 2a *If $(j, j')\in\mathcal{J}$ then*

$$P\left[\left(c_j(\mathbf{x})=1\right)\cap\left(c_{j'}(\mathbf{x})=1\right)\right]=P\left[c_j(\mathbf{x})=1\right]+P\left[c_{j'}(\mathbf{x})=1\right]-1+\epsilon_{jj'} \tag{51}$$

for an input–output vector $\mathbf{x}\in V$ selected uniformly at random.

The quantity $\epsilon_{jj'}$ is a measure of the "overlap error". We can use it to prove relaxed versions of Lemma 1 and Theorem 1.

Lemma 1a *Assume Condition 1 and $(j, j') \in \mathcal{J}$. Then the sum of the correctness values of the corresponding weak classifiers is nonnegative everywhere with probability $1 - \epsilon_{jj'}$, namely*

$$P\left[c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \geq 0\right] = 1 - \epsilon_{jj'} \tag{52}$$

for an input–output vector $\mathbf{x} \in \mathcal{V}$ selected uniformly at random.

Proof The proof closely mimics that of Lemma 1.

$$\begin{aligned}
&P\left[(c_j(\mathbf{x}) = 1) \cup (c_{j'}(\mathbf{x}) = 1)\right] \\
&= P\left[c_j(\mathbf{x}) = 1\right] + P\left[c_{j'}(\mathbf{x}) = 1\right] - P\left[(c_j(\mathbf{x}) = 1) \cap (c_{j'}(\mathbf{x}) = 1)\right] \\
&= P\left[c_j(\mathbf{x}) = 1\right] + P\left[c_{j'}(\mathbf{x}) = 1\right] - P\left[c_j(\mathbf{x}) = 1\right] - P\left[c_{j'}(\mathbf{x}) = 1\right] + 1 - \epsilon_{jj'} \\
&= 1 - \epsilon_{jj'}
\end{aligned} \tag{53}$$

by Condition 2a. As in the proof of Lemma 1, this implies

$$P\left[c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \in \{0, 2\}\right] = 1 - \epsilon_{jj'}. \tag{54}$$

We can now replace Theorem 1 by a lower bound on the success probability when Condition 2 is replaced by the weaker Condition 2a. Let us first define an imperfect strong classifier as follows:

Definition 7 A strong classifier is ϵ-perfect if, for $\mathbf{x} \in \mathcal{V}$ chosen uniformly at random, it correctly classifies $\mathbf{x}$ [i.e., $Q_{\mathbf{w}}(\mathbf{x}) = y(\mathbf{x})$] with probability at least $1 - \epsilon$.

Theorem 3 *A strong classifier comprised solely of a set of weak classifiers satisfying Conditions 1, 2a and 3 is ϵ-perfect, where $\epsilon = \sum_{(j,j') \in \mathcal{J}} \epsilon_{jj'}$.*

Proof It suffices to show that the correctness sum is positive with probability 1 minus the sum of the overlap errors when Conditions 1, 2a and 3 are satisfied, namely

$$P\left[\sum_{(j,j') \in \mathcal{J}} c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) > 0\right] \geq 1 - \sum_{(j,j') \in \mathcal{J}} \epsilon_{jj'}. \tag{55}$$

Now, by definition $c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \in \{-2, 0, 2\}$, and the correctness sum of at least one of the pairs must be negative in order for the correctness sum over all weak classifiers in $\mathcal{J}$ to be negative, so that

$$P\left[\sum_{(j,j') \in \mathcal{J}} c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) < 0\right] \tag{56a}$$

$$\leq P\left[\exists (j, j') \in \mathcal{J} : c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) = -2\right]. \tag{56b}$$

However, we also need to exclude the case of all weak classifier pairs summing to zero (otherwise the strong classifier can be inconclusive). This case is partially excluded by virtue of Condition 3, which tells us that $\mathcal{V}$ as a whole is always covered by a correct majority vote. Formally,

$$
P\left[\sum_{(j,j')\in\mathcal{J}} c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) = 0\right] = P\left[\bigcap_{(j,j')\in\mathcal{J}} \left(c_j(\mathbf{x}) + c_{j'}(\mathbf{x})\right) = 0\right]
$$
$$
= 1 - P\left[\exists(j,j') \in \mathcal{J} : c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) > 0\right]
$$
$$
= 0, \tag{57}
$$

where in the last equality we invoked the calculation leading from Eqs. (46c) to (47), which only required Condition 3. Alternatively, we could use Condition 3a to prove that $P\left[\sum_{(j,j')\in\mathcal{J}} c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) + c_a(\mathbf{x}) = 0\right] = 0$. There is another way for the classifier to return an inconclusive result: if one weak classifier pair has a correctness sum of 2 and another weak classifier pair has a correctness sum of -2. This case is included in the bound in Eq. (56b) because one of the weak classifier pairs in this scenario has a negative correctness sum. We can thus conclude that the strict inequality in Eq. (56a) can be replaced by $\leq$.

Now, the probability of there being one weak classifier pair such as in Eq. (56b) cannot be greater than the probability of at least one of the pairs having a negative correctness sum, which in turn—by the union bound—cannot be greater than the sum of such probabilities:

$$
\text{Eq. (56b)} \leq P\left[\bigcup_{(j,j')\in\mathcal{J}} \left(c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) = -2\right)\right]
$$
$$
\leq \sum_{(j,j')\in\mathcal{J}} P\left[c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) = -2\right]
$$
$$
= \sum_{(j,j')\in\mathcal{J}} \epsilon_{jj'}, \tag{58}
$$

where the last equality follows from Lemma 1a. This proves Eq. (55).

It is interesting to note that—as alluded to in this proof—if we were to drop Conditions 3 and 3a, then Eq. (55) would become

$$
P\left[\sum_{(j,j')\in\mathcal{J}} c_j(\mathbf{x}) + c_{j'}(\mathbf{x}) \geq 0\right] \geq 1 - \sum_{(j,j')\in\mathcal{J}} \epsilon_{jj'}
$$

(note the change from > to $\geq$), so that Theorem 3 would change to a statement about inconclusive ϵ-perfect strong classifiers, which can—with finite probability—yield a "don't-know" answer. This may be a useful tradeoff if it turns out to be difficult to construct a set of weak classifiers satisfying Condition 3 or 3a.

4.4 An alternate weight optimization problem

The conditions and results established in the previous subsection for correctness of the strong classifier suggest the creation of an alternate weight optimization problem to select the weak classifiers that will be included in the final majority vote, replacing the optimization problem of Sect. 3.4. The new optimization problem is defined over the space of *pairs* of weak classifiers, rather than singles, which can be constructed using elements of the set $\mathcal{A} \times \mathcal{A}$, with $\mathcal{A}$ as defined in Condition 1. We define the *ideal pair weight* as

$$\tilde{w}_{ij} = \begin{cases} 1 & (i, j) \in \mathcal{J} \times \mathcal{J} \\ 0 & \text{otherwise} \end{cases}, \tag{59}$$

Since we do not know the set $\mathcal{J}$ *a priori*, we shall define a QUBO whose solutions $w_{ij} \in \{0, 1\}$, with $(i, j) \in \mathcal{A} \times \mathcal{A}$, will be an approximation to the ideal pair weights $\tilde{w}_{ij}$. In the process, we shall map the pair weight bits w_{ij} to qubits. Each w_{ij} determines whether its corresponding pair of weak classifiers, h_i and h_j, will be included in the new strong classifier, which can thus be written as:

$$\begin{aligned} Q^{\text{pair}}(\mathbf{x}) &= \text{sign}\left[R_{\mathbf{w}^{\text{pair}}}(\mathbf{x}) \right] \\ &= \text{sign}\left[\sum_{(i,j) \in \mathcal{A} \times \mathcal{A}} w_{ij} \left(h_i(\mathbf{x}) + h_j(\mathbf{x}) \right) \right] \end{aligned} \tag{60}$$

Recall that we do not know the w_{ij} *a priori*; they are found in our approach via the solution of a QUBO, which we set up as follows:

$$\mathbf{w}_{\text{pair}}^{\text{opt}} = \arg \min_{\mathbf{w}} \left[\sum_{(i,j) \in \mathcal{A} \times \mathcal{A}} \alpha_{ij} w_{ij} + \sum_{(i,j) \neq (k,l) \in \mathcal{A} \times \mathcal{A}} J_{ijkl} w_{ij} w_{kl} \right], \tag{61}$$

where the second term is a double sum over all sets of unequal pairs. The solution of this QUBO will provide us with an approximation to the set $\mathcal{J}$, which yields the desired set of weak classifiers as in Eq. (41). Sparsity can be enforced as in Eq. (22) by replacing α_{ij} with $\alpha_{ij} + \lambda$, where $\lambda > 0$, i.e., by including a penalty proportional to $\|\mathbf{w}\|_0$.

The terms α_{ij} and J_{ijkl} reward compliance with Conditions 2 and 3, respectively. To define α_{ij}, we first define the modified correctness function $c_i' : \mathcal{T} \mapsto \{0, 1\}$, where $\mathcal{T}$ is the training set (13):

$$c_i'(\mathbf{x}_s, y_s) = \frac{1}{2}\left(h_i(\mathbf{x}_s) y_s + 1 \right) = \begin{cases} 1 & h_i(\mathbf{x}_s) = y_s \\ 0 & h_i(\mathbf{x}_s) \neq y_s \end{cases} \tag{62}$$

Below we write $c_i'(s)$ in place of $c_i'(\mathbf{x}_s, y_s)$ for notational simplicity. The term α_{ij} rewards the pairing of weak classifiers which classify the minimal number of vectors $\mathbf{x}$ incorrectly at the same time, as specified by Condition 2. Each pair included gains negative weight for the training set vectors its members classify correctly, but is also given

a positive penalty for any vectors classified incorrectly by both weak classifiers at once:

$$\alpha_{ij} = -\frac{1}{|\mathcal{T}|} \sum_{s=1}^{|\mathcal{T}|} \left[c_i'(s) + c_j'(s) - \left(1 - c_i'(s)\right)\left(1 - c_j'(s)\right) \right] \tag{63}$$

The term J_{ijkl} penalizes the inclusion of pairs that are too similar to each other, as codified in Condition 3. This is accomplished by assigning a positive weight for each vector that is classified correctly by two pairs at once:

$$J_{ijkl} = \frac{1}{|\mathcal{T}|} \sum_{s=1}^{|\mathcal{T}|} c_i'(s) c_j'(s) c_k'(s) c_l'(s) \tag{64}$$

We now have a QUBO for the alternate weight optimization problem. This can be translated to the Ising Hamiltonian as with the original optimization problem in Sect. 3.5. We again map from our QUBO variables w_{ij} to variables $q_{ij} = 2(w_{ij} - 1/2)$, yielding the following optimization function:

$$\mathbf{q}_{\text{pair}}^{\text{opt}} = \arg\min_{\mathbf{q}} \left[\frac{1}{2} \sum_{(i,j)\in\mathcal{A}\times\mathcal{A}} \beta_{ij} q_{ij} + \frac{1}{4} \sum_{(i,j)\neq(k,l)\in\mathcal{A}\times\mathcal{A}} J_{ijkl} q_{ij} q_{kl} \right], \tag{65}$$

where

$$\beta_{ij} = \alpha_{ij} + \frac{1}{2} \left(\sum_{(k,l)\in\mathcal{A}\times\mathcal{A}; (k,l)\neq(i,j)} J_{ijkl} + J_{klij} \right). \tag{66}$$

Constant terms were omitted because they have no bearing on the minimization. This optimization function is now suitable for direct translation to the final Hamiltonian for an AQC:

$$H_F = \frac{1}{2} \sum_{(i,j)\in\mathcal{A}\times\mathcal{A}} \beta_{ij} Z_{ij} + \frac{1}{4} \sum_{(i,j)\neq(k,l)\in\mathcal{A}\times\mathcal{A}} J_{ijkl} Z_{ij} Z_{kl}. \tag{67}$$

The qubits now represent weights on pairs rather than on an individual classifier. Z_{ij} is therefore the Pauli σ_z operator on the qubit assigned to the pair $(i, j) \in \mathcal{A}\times\mathcal{A}$. Using $|\mathcal{A}|^2$ qubits, this approach will give the optimal combination of weak classifier pairs over the training set according to the conditions set forth previously.

5 Using strong classifiers in quantum-parallel

Now let us suppose that we have already trained our strong classifier and found the optimal weight vector $\mathbf{w}^{\text{opt}}$ or $\mathbf{w}_{\text{pair}}^{\text{opt}}$. For simplicity we shall henceforth limit our discussion to $\mathbf{w}^{\text{opt}}$. We can use the trained classifier to classify new input–output pairs $\mathbf{x} \notin \mathcal{T}$ to

decide whether they are correct or erroneous. In this section we shall address the question of how we can further obtain a quantum speedup in exhaustively testing *all* exponentially many ($2^{N_{in}+N_{out}}$) input–output pairs $\mathbf{x}$. The key observation in this regard is that if we can formulate software error testing as a minimization problem over the space $\mathcal{V}$ of all input–output pairs $\mathbf{x}$, then an AQC algorithm will indeed perform a quantum-parallel search over this entire space, returning as the ground state an erroneous state.

5.1 Using two strong binary classifiers to detect errors

Recall that we are concerned with the detection of vectors $\mathbf{x} \in \mathcal{V}$ that are erroneous and implemented (Eq. 7). To accomplish this, we use two strong classifiers. The *specification classifier* is the binary classifier developed in Sect. 3. Ideally, it behaves as follows:

$$Q_{\mathbf{w}}(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \in \hat{S} \\ -1 & \mathbf{x} \notin \hat{S} \end{cases} \tag{68}$$

The second classifier, which we will call the *implementation classifier*, determines whether or not an input–output vector is in the program as implemented. It is constructed in the same way as $Q_{\mathbf{w}}(\mathbf{x})$, but with its own appropriate training set. Ideally, it behaves as follows:

$$T_{\mathbf{z}}(\mathbf{x}) = \begin{cases} 1 & \mathbf{x} \notin S \\ -1 & \mathbf{x} \in S \end{cases} \tag{69}$$

The four possible combinations represented by Eqs. (68) and (69) correspond to the four cases covered by Definitions 1–4. The worrisome input–output vectors, those that are erroneous and implemented, cause both classifiers to evaluate to -1.

5.2 Formal criterion

As a first step, suppose we use the optimal weights vector in the original strong specification classifier. We then have, from (15),

$$Q^{opt}(\mathbf{x}) = \text{sign}\left[R_{\mathbf{w}^{opt}}(\mathbf{x})\right] = \text{sign}\left[\sum_{i=1}^{N} w_i^{opt} h_i(\mathbf{x})\right] \tag{70}$$

This, of course, is imprecise since our adiabatic algorithm solves a relaxed optimization problem (i.e., returns $\mathbf{w}^{opt}$, not $\mathbf{w}'^{opt}$), but we shall assume that the replacement is sufficiently close to the true optimum for our purposes. With this caveat, Eq. (70) is the optimal strong specification classifier for a given input–output vector $\mathbf{x}$, with the classification of $\mathbf{x}$ as erroneous if $Q^{opt}(\mathbf{x}) = -1$ or as correct if $Q^{opt}(\mathbf{x}) = +1$.

The strong implementation classifier is constructed similarly to the specification classifier:

$$T^{\text{opt}}(\mathbf{x}) = \text{sign}\left[U_{\mathbf{z}^{\text{opt}}}(\mathbf{x})\right] = \text{sign}\left[\sum_{i=1}^{N} z_i^{\text{opt}} h_i(\mathbf{x})\right] \tag{71}$$

Here, h_i are the same weak classifiers as those used to train the specification classifier, but T^{opt} is constructed independently from a training set $\mathcal{T}'$ which may or may not overlap with $\mathcal{T}$. This training set is labeled according to the possibility or impossibility of producing the input–output pairs in $\mathcal{T}'$ from the implemented program. The result of this optimization is the weight vector $\mathbf{z}^{\text{opt}}$.

Given the results of the classifiers $Q^{\text{opt}}(\mathbf{x})$ and $T^{\text{opt}}(\mathbf{x})$ for any vector $\mathbf{x}$, the V&V task of identifying whether or not $\mathbf{x} \in (S \cap \neg\hat{S})$ reduces to the following. Any vector $\mathbf{x}$ is flagged as erroneous and implemented if $Q^{\text{opt}}(\mathbf{x}) + T^{\text{opt}}(\mathbf{x}) = -2$. We stress once more that, due to our use of the relaxed optimization to solve for $\mathbf{w}^{\text{opt}}$ and $\mathbf{z}^{\text{opt}}$, a flagged $\mathbf{x}$ may in fact be neither erroneous nor implemented, i.e., our procedure is susceptible to both false positives and false negatives.

5.3 Relaxed criterion

As was the case with Eq. (18), $Q^{\text{opt}} + T^{\text{opt}}$ is unfortunately not directly implementable in AQC, but a simple relaxation is. The trick is again to remove the sign function, this time from (70) and (71), and consider the sum of the two classifiers' majority vote functions directly as an energy function:

$$C^{\text{opt}}(\mathbf{x}) = R_{\mathbf{w}^{\text{opt}}}(\mathbf{x}) + U_{\mathbf{z}^{\text{opt}}}(\mathbf{x}) \tag{72}$$

The combination of the two classifiers gives different results for vectors falling under each of the Definitions from Sect. 2.2.2.

Case 1: $\mathbf{x} \notin \hat{S}$ and $\mathbf{x} \in S$

The vector $\mathbf{x}$ is an error implemented in the program and manifests a software error. These vectors gain negative weight from both classifiers $R_{\mathbf{w}^{\text{opt}}}$ and $U_{\mathbf{z}^{\text{opt}}}$. Vectors falling under this definition should receive the lowest values of C^{opt}, if any such vectors exist.

Case 2: $\mathbf{x} \in \hat{S}$ and $\mathbf{x} \in S$

The vector $\mathbf{x}$ satisfies the don't-worry condition, that is, it is a correct input–output string, part of the ideal program $\hat{P}$. In this case, $R_{\mathbf{w}^{\text{opt}}} > 0$ and $U_{\mathbf{z}^{\text{opt}}} < 0$. In the programs quantum V&V is likely to be used for, with very infrequent, elusive errors, the specification and implementation will be similar and the negative weight of $U_{\mathbf{z}^{\text{opt}}} < 0$ should be moderated enough by the positive influence of $R_{\mathbf{w}^{\text{opt}}} > 0$ that don't-worry vectors should not populate the lowest-lying states.

Case 3: $\mathbf{x} \in \hat{S}$ and $\mathbf{x} \notin S$

The input portion of the vector $\mathbf{x}$ is a don't-care condition. It does not violate any program specifications, but is not important enough to be specifically addressed in the implementation. This vector will gain positive weight from both $R_{\mathbf{w}^{\text{opt}}}$ and $U_{\mathbf{z}^{\text{opt}}}$ and should therefore never be misidentified as an error.

Case 4: $\mathbf{x} \notin \hat{S}$ and $\mathbf{x} \notin S$

The vectors $\mathbf{x}$ in this category would be seen as erroneous by the program specification - if they ever occurred. Because they fall outside the program implementation S, they are not the errors we are trying to find. This case is similar to the don't-worry situation in that the two strong classifiers will have opposite signs, in this case $R_{\mathbf{w}^{\mathrm{opt}}} < 0$ and $U_{\mathbf{z}^{\mathrm{opt}}} > 0$. By the same argument as Definitions 2 and 4 vectors should not receive more negative values of C^{opt} than the targeted errors.

Having examined the values of $C^{\mathrm{opt}}(\mathbf{x})$ for the relevant categories of $\mathbf{x}$, we can formulate error detection as the following minimization problem:

$$\mathbf{x}_e = \arg\min_{\mathbf{x}} C^{\mathrm{opt}}(\mathbf{x}). \tag{73}$$

Suppose the algorithm returns a solution $\mathbf{x}_e$ (e for "error"). We then need to test that it is indeed an error, which amounts to checking that it behaves incorrectly when considered as an input–output pair in the program implementation P. Note that testing that $R_{\mathbf{w}^{\mathrm{opt}}}(\vec{x}_e) < 0$ is insufficient, since our procedure involved a sequence of relaxations.

5.4 Adiabatic implementation of the relaxed criterion

In order to implement the error identification strategy (73) we need to consider

$$C^{\mathrm{opt}}(\mathbf{x}) = \sum_{i=1}^{N} (w_i^{\mathrm{opt}} + z_i^{\mathrm{opt}}) h_i(\mathbf{x}) \tag{74}$$

as an energy function. We then consider $C^{\mathrm{opt}}(\vec{x})$ as the final Hamiltonian H_F for an AQC, with Hilbert space spanned by the basis $\{|\mathbf{x}\rangle\}$. The AQC will then find the state which minimizes $C^{\mathrm{opt}}(\mathbf{x})$ out of all $2^{N_{\mathrm{in}}+N_{\mathrm{out}}}$ basis states and thus identify an error candidate. Because the AQC always returns *some* error candidate, our procedure never generates false negatives. However, Cases 2 and 4 would correspond to false positives, if an input–output vector satisfying either one of these cases is found as the AQC output.

We can rely on the fact that the AQC actually returns a (close approximation to the) Boltzmann distribution

$$\Pr[\mathbf{x}] = \frac{1}{Z} \exp[-C^{\mathrm{opt}}(\mathbf{x})/(k_B T)], \tag{75}$$

where k_B is the Boltzmann constant, T is the temperature, and

$$Z = \sum_{\mathbf{x}} \exp[-C^{\mathrm{opt}}(\mathbf{x})/(k_B T)] \tag{76}$$

is the partition function. For sufficiently low temperature this probability distribution is sharply peaked around the ground state, with contributions from the first few excited states. Thus we can expect that even if there is a low-lying state that has been pushed there by only one of the two binary classifiers Q^{opt} or T^{opt}, the AQC

will return a nearby state which is both erroneous and implemented some of the time and an error will still be detected. Even if the undesirable state [$\mathbf{x} \in \hat{S}$ and $\mathbf{x} \in S$, or $\mathbf{x} \notin \hat{S}$ and $\mathbf{x} \notin S$] is the ground state, and hence all erroneous states [$\mathbf{x} \notin \hat{S}$ and $\mathbf{x} \in S$] are excited states, their lowest energy member will be found with a probability that is $e^{-\Delta(t_F)/(k_B T)}$ smaller than the unlooked-for state, where $\Delta(t_F)$ is the energy gap to the first excited state at the end of the computation. Provided $k_B T$ and $\Delta(t_F)$ are of the same order, this probability will be appreciable.

To ensure that errors which are members of the training set are never identified as ground states we construct the training set $\mathcal{T}$ so that it only includes correct states, i.e., $y_s = +1 \, \forall s$. This has the potential drawback that the classifier never trains directly on errors. It is in principle possible to include errors in the training set ($y_s = -1$) by adding another penalty term to the strong classifier which directly penalizes such training set members, but whether this can be done without introducing many-body interactions in H_F is a problem that is beyond the scope of this work.

Also beyond the scope of this work is an analysis of the gap for the quantum testing step of the algorithm. The problems are simply too large to solve analytically, and too dependent on the outcome of the learning step to be easily attacked numerically. However, verification and validation of software is such an important problem that quantum formulations must be explored. Even if the scaling of the algorithm proves unfavorable for the V&V problem, other anomaly detection applications may exhibit better scaling due to different outcomes from the learning step.

5.5 Choosing the weak classifiers

Written in the form $\sum_{i=1}^{N} (w_i^{\text{opt}} + z_i^{\text{opt}}) h_i(\mathbf{x})$, the energy function $C^{\text{opt}}(\mathbf{x})$ is too general, since we haven't yet specified the weak classifiers $h_i(\mathbf{x})$. However, we are free to choose these so as to mold $C^{\text{opt}}(\mathbf{x})$ into a Hamiltonian that is physically implementable in AQC.

Suppose, e.g., that $h_i(\mathbf{x})$ measures a Boolean relationship defined by a function $f_i : \{0, 1\}^{\ell} \mapsto \{0, 1\}$ between several bits of the input–output vector; $x_k = \text{bit}_k(\mathbf{x})$, the kth bit of $\mathbf{x} \in \mathcal{V}$. For example,

$$h_i(\mathbf{x}) = (x_{i_3} == f_i(x_{i_1}, x_{i_2})), \tag{77}$$

where "$a == b$" evaluates to 1 if $a = b$ or to 0 if $a \neq b$. Here i_1 and i_2 are the positions of two bits from the input vector $\vec{x}_{\text{in}}$ and i_3 is the position of a bit from the output vector $\mathbf{x}_{\text{out}}$, so that h_i measures a correlation between inputs and outputs. The choice of this particular form for the weak classifiers is physically motivated, as it corresponds to at most three-body interactions between qubits, which can all be reduced to two-body interaction by the addition of ancilla qubits (see below). Let us enumerate

these weak classifiers. The number of different Boolean functions f_i is 2^{2^ℓ} [49].[4] Much more efficient representations are possible under reasonable assumptions [50], but for the time being we shall not concern ourselves with these. In the example of the classifier (77) there are $N_{\mathrm{in}}(N_{\mathrm{in}} - 1)$ input bit combinations for each of the N_{out} output bits. The number of different Boolean functions in this example, where $\ell = 2$, is $2^{2^2} = 16$. Thus the dimension of the "dictionary" of weak classifiers is

$$N = 16 N_{\mathrm{in}}(N_{\mathrm{in}} - 1) N_{\mathrm{out}} \tag{78}$$

for the case of Eq. (77).

We wish to find a two-local quantum implementation for each $h_i(\mathbf{x})$ in the dictionary. It is possible to find a two-local implementation for any three-local Hamiltonian using so-called "perturbation gadgets", or three ancilla bits for each three-local term included [51], but rather than using the general method we rely on a special case which will allow us to use only one ancilla bit per three-local term. We first devise an *intermediate form* function using products of the same bits $x_i \in \{0, 1\}$ used to define the logical behavior of each weak classifier. This function will have a value of 1 when the Boolean relationship specified for $h_i(\mathbf{x})$ is true, and -1 otherwise. For example, consider function number 8, $x_{i_3} == x_{i_1} \wedge x_{i_2}$, the AND function. Its intermediate form is $4 x_{i_1} x_{i_2} x_{i_3} - 2 \left(x_{i_3} + x_{i_1} x_{i_2} \right) + 1$. For the bit values $(x_{i_1}, x_{i_2}, x_{i_3}) = (0, 0, 0)$, the value of the intermediate function is 1, and the Boolean form is true: 0 AND 0 yields 0. If instead we had the bit values $(x_{i_1}, x_{i_2}, x_{i_3}) = (0, 0, 1)$, the intermediate form would yield -1, and the Boolean form would be false, because the value for x_{i_3} does not follow from the values for x_{i_1} and x_{i_2}.

The two-body implementation form is obtained in two steps from the intermediate form. First, an ancilla bit tied to the product of the two input bits, $x_a = x_{i_1} x_{i_2}$, is substituted into any intermediate form expressions involving three-bit products. This is permissible because such an ancilla can indeed be created by introducing a penalty into the final Hamiltonian for any states in which the ancilla bit is not equal to the product $x_{i_1} x_{i_2}$. We detail this method below. Then, the modified intermediate expression is translated into a form that uses bits valued as $x_i' \in \{-1, 1\}$ rather than $x_i \in \{0, 1\}$ using the equivalence $x_i = 2 x_i' - 1$. The modified intermediate form is now amenable to using the implemented qubits. Note that the Pauli matrix Z_i acts on a basis ket $|\mathbf{x}\rangle$ as

$$Z_i |\mathbf{x}\rangle = (-1)^{\mathrm{bit}_i(\mathbf{x})} |\mathbf{x}\rangle. \tag{79}$$

This means that we can substitute Z_i for x_i' and $Z_i \otimes Z_j$ for $x_i' x_j'$ in the intermediate form, resulting in the implementation form given in Column 4 of Table 1. Some weak classifiers do not involve three-bit interactions. Their implementation forms were devised directly, a simple process when there is no need for inclusion of an ancilla.

---

[4] Any Boolean function of ℓ variables can be uniquely expanded in the form $f_i(x_1, \ldots, x_\ell) = \sum_{\alpha=0}^{2^\ell - 1} \epsilon_{i\alpha} s_\alpha$, where $\epsilon_{i\alpha} \in \{0, 1\}$ and s_α are the 2^ℓ "simple" Boolean functions $s_0 = x_1 x_2 \cdots x_\ell$, $s_1 = x_1 x_2 \cdots \overline{x_\ell}$, ..., $s_{2^\ell - 1} = \overline{x_1}\, \overline{x_2} \cdots \overline{x_\ell}$, where $\overline{x}$ denotes the negation of the bit x. Since each $\epsilon_{i\alpha}$ can assume one of two values, there are 2^{2^ℓ} different Boolean functions.

Table 1 All 16 Boolean functions f_i of two binary variables, and their implementation form in terms of the Pauli matrices Z_{i_j} acting on single qubits or pairs of qubits $j \in \{1, 2, 3\}$

Function #	Boolean logic	Intermediate form	Implementation form
$i = 0$	$x_{i_3} == 0$	Not applicable	$-Z_{i_3}$
$i = 1$	$x_{i_3} == \overline{(x_{i_1} \vee x_{i_2})}$	$4\left(x_{i_1}x_{i_2}x_{i_3} - x_{i_1}x_{i_3} - x_{i_2}x_{i_3}\right) - 2\left(x_{i_1}x_{i_2} - x_{i_1} - x_{i_2} - x_{i_3}\right) - 1$	$Z_a \otimes Z_{i_3} - Z_{i_1} \otimes Z_{i_3} - Z_{i_2} \otimes Z_{i_3}$
$i = 2$	$x_{i_3} == \overline{x_{i_1}} \wedge x_{i_2}$	$4(-x_{i_1}x_{i_2}x_{i_3} + x_{i_2}x_{i_3}) + 2(-x_{i_3} + x_{i_1}x_{i_2} - x_{i_2}) + 1$	$-Z_a \otimes Z_{i_3} + Z_{i_2} \otimes Z_{i_3} - Z_{i_3}$
$i = 3$	$x_{i_3} == \overline{x_{i_1}}$	Not applicable	$-Z_{i_3} \otimes Z_{i_1}$
$i = 4$	$x_{i_3} == x_{i_1} \wedge \overline{x_{i_2}}$	$4(x_{i_1}x_{i_3} - x_{i_1}x_{i_2}x_{i_3}) - 2(x_{i_1} - x_{i_1}x_{i_2} + x_{i_3}) + 1$	$Z_{i_1} \otimes Z_{i_3} - Z_a \otimes Z_{i_3} - Z_{i_3}$
$i = 5$	$x_{i_3} == \overline{x_{i_2}}$	Not applicable	$-Z_{i_3} \otimes Z_{i_2}$
$i = 6$	$x_{i_3} == x_{i_1} \oplus x_{i_2}$	$-8x_{i_1}x_{i_2}x_{i_3} + 4(x_{i_1}x_{i_3} + x_{i_2}x_{i_3} + x_{i_1}x_{i_2}) - 2(x_{i_1} + x_{i_2} + x_{i_3}) + 1$	$-2Z_a \otimes Z_{i_3} + Z_{i_1} \otimes Z_{i_3} + Z_{i_2} \otimes Z_{i_3} - Z_{i_3}$
$i = 7$	$x_{i_3} == \overline{(x_{i_1} \wedge x_{i_2})}$	$-4x_{i_1}x_{i_2}x_{i_3} + 2(x_{i_3} + x_{i_1}x_{i_2}) - 1$	$-Z_a \otimes Z_{i_3}$
$i = 8$	$x_{i_3} == x_{i_1} \wedge x_{i_2}$	$4x_{i_1}x_{i_2}x_{i_3} - 2(x_{i_3} + x_{i_1}x_{i_2}) + 1$	$Z_a \otimes Z_{i_3}$
$i = 9$	$x_{i_3} == \overline{(x_{i_1} \oplus x_{i_2})}$	$8x_{i_1}x_{i_2}x_{i_3} - 4(x_{i_1}x_{i_3} + x_{i_2}x_{i_3} + x_{i_1}x_{i_2}) + 2(x_{i_1} + x_{i_2} + x_{i_3}) - 1$	$2Z_a \otimes Z_{i_3} - Z_{i_1} \otimes Z_{i_3} - Z_{i_2} \otimes Z_{i_3} + Z_{i_3}$
$i = 10$	$x_{i_3} == x_{i_2}$	Not applicable	$Z_{i_3} \otimes Z_{i_2}$
$i = 11$	$x_{i_3} == \overline{x_{i_1}} \vee x_{i_2}$	$-4(x_{i_1}x_{i_3} - x_{i_1}x_{i_2}x_{i_3}) + 2(x_{i_1} - x_{i_1}x_{i_2} + x_{i_3}) - 1$	$-Z_{i_1} \otimes Z_{i_3} + Z_a \otimes Z_{i_3} + Z_{i_3}$
$i = 12$	$x_{i_3} == x_{i_1}$	Not applicable	$Z_{i_3} \otimes Z_{i_1}$
$i = 13$	$x_{i_3} == x_{i_1} \vee \overline{x_{i_2}}$	$-4(-x_{i_1}x_{i_2}x_{i_3} + x_{i_2}x_{i_3}) - 2(-x_{i_3} + x_{i_1}x_{i_2} - x_{i_2}) - 1$	$Z_a \otimes Z_{i_3} - Z_{i_2} \otimes Z_{i_3} + Z_{i_3}$
$i = 14$	$x_{i_3} == x_{i_1} \vee x_{i_2}$	$-4(x_{i_1}x_{i_2}x_{i_3} - x_{i_1}x_{i_3} - x_{i_2}x_{i_3}) + 2(x_{i_1}x_{i_2} - x_{i_1} - x_{i_2} - x_{i_3}) + 1$	$-Z_a \otimes Z_{i_3} + Z_{i_1} \otimes Z_{i_3} + Z_{i_2} \otimes Z_{i_3}$
$i = 15$	$x_{i_3} == 1$	Not applicable	Z_{i_3}

The subscript a in the implementation form column denotes an ancilla qubit, tied to qubits i_1 and i_2 via $x_a = x_{i_1}x_{i_2}$, used to reduce all qubit interactions to at most two-body

We have reduced the dictionary functions from three-bit to two-bit interactions by adding an ancilla bit to represent the product of the two input bits involved in the function. Therefore, the maximum number of qubits needed to implement this set of weak classifiers on a quantum processor is $Q = N_{in} + N_{out} + N_{in}^2$. In practice, it is likely to be significantly less because not every three-bit correlation will be relevant to a given classification problem.

Let us now discuss how the penalty function is introduced. For example, consider again the implementation of weak classifier function $i = 8$, whose intermediate form involves three-qubit products, which we reduced to two-qubit interactions by including x_a.

We ensure that x_a does indeed represent the product it is intended to by making the function a sum of two terms: the product of the ancilla qubit and the remaining qubit from the original product, and a term that adds a penalty if the ancilla is not in fact equal to the product of the two qubits it is meant to represent, in this case $f_{penalty} = x_{i_1} x_{i_2} - 2(x_{i_1} + x_{i_2})x_a + 3x_a$. In the case where $(x_{i_1}, x_{i_2}, x_a) = (1, 0, 0)$, $f_{penalty} = 0$, but in a case where x_a does not represent the intended product such as $(x_{i_1}, x_{i_2}, x_a) = (1, 0, 1)$, $f_{penalty} = 1$. In fact, the penalty function behaves as follows:

$$f_{penalty} = \begin{cases} 0 & x_a = x_{i_1} x_{i_2} \\ positive & otherwise \end{cases} \tag{80}$$

In the end, we have the modified intermediate form $f_8 = 4x_a x_{i_3} - 2\left(x_{i_3} + x_a\right) + 1 + f_{penalty}$, which involves only two-qubit interactions. This would be implemented on the quantum computer as the sum of two Hamiltonian terms:

$$H_8 = Z_a \otimes Z_{i_3}, \tag{81}$$

from the implementation column of Table 1, and

$$H_{penalty}(i_1, i_2) = \frac{1}{4}Z_{i_1} \otimes Z_{i_2} - \frac{1}{2}Z_{i_1} \otimes Z_a - \frac{1}{2}Z_{i_2} \otimes Z_a$$
$$-\frac{1}{4}Z_{i_1} - \frac{1}{4}Z_{i_2} + \frac{1}{2}Z_a + \frac{3}{4}, \tag{82}$$

the implementation form of $f_{penalty}$, so a Hamiltonian to find input–output vectors classified negatively by this weak classifier would be

$$H_{weak} = H_8 + H_{penalty}(i_1, i_2). \tag{83}$$

When the strong classifier is implemented as a whole, multiple weak classifiers with weight 1 may use the same two input bits, and therefore share an ancilla bit that is the product of those input bits. When this is the case, it is sufficient to add the penalty function to the final Hamiltonian once, though the ancilla is used multiple times.

The inclusion of ancilla qubits tied to products of other qubits and their associated penalties need not interfere with the solution of the V&V problem, although the ancilla penalty terms must appear in the same final Hamiltonian as this optimization. If the

ancilla penalty terms are made reasonably large, they will put any states in which the ancillas do not represent their intended products (states which are in fact outside of $\mathcal{V}$) far above the levels at which errors are found. For instance, consider an efficient, nearly optimal strong classifier closely approximating the conditions set forth in Sect. 4. Such a classifier makes its decision on the strength of two simultaneously true votes. If two such classifiers are added together, as in the verification problem, the lowest energy levels will have an energy near -4. If the penalty on a forbidden ancilla state is more than a reasonable 4 units, such a state should be well clear of the region where errors are found.

This varied yet correlation-limited set of weak classifiers fits nicely with the idea of tracking intermediate spaces (Eq. 11), where we can use an intermediate space $\mathcal{I}_j$ to construct a set of weak classifiers feeding into the next intermediate space $\mathcal{I}_{j+1}$. This is further related to an obvious objection to the above classifiers, which is that they ignore any correlations involving four or more bits, without one-, two-, or three-bit correlations. By building a hierarchy of weak classifiers, for intermediate spaces, such correlations can hopefully be accounted for as they build up by keeping track instead of one-, two-, and three-bit terms as the program runs.

5.6 QUBO-AQC quantum parallel testing

With the choice of Boolean functions for the weak classifiers, the quantum implementation of the energy function $C^{\text{opt}}(\mathbf{x})$ (Eq. 74) becomes

$$H_F^{\text{test}} = \sum_{i=1}^{N} (w_i^{\text{opt}} + z_i^{\text{opt}}) H_i + \sum_{j \neq k} H_{\text{penalty}}(j, k), \tag{84}$$

where H_i denotes the implemented form given in the third column of Table 1, and the indices $j, k \in \{1, \ldots, N_{\text{in}}\}$ denote all possible pairings of input qubits tied to ancillas. The ground state of H_F^{test}, which corresponds to the optimal weight sets w_i^{opt} and z_i^{opt} derived from the set of weak classifiers detailed in Sect. 5.5, is an erroneous state, which, by construction, is not a member of the training set $\mathcal{T}$.

How do we construct the AQC such that all input–output pairs $\mathbf{x}$ are tested in parallel? This is a consequence of the adiabatic interpolation Hamiltonian (26), and in particular the initial Hamiltonian H_I of the type given in Eq. (27). The ground state of this positive semi-definite H_I is an equal superposition over all input–output vectors, i.e., $H_I \sum_{\mathbf{x} \in \mathcal{V}} |\mathbf{x}\rangle = 0$, and hence when we implement the AQC every possible $\mathbf{x}$ starts out as a candidate for the ground state. The final (Boltzmann) distribution of observed states strongly favors the manifold of low energy states, and by design these will be implemented erroneous states, if they exist.

6 Sample problem implementation

In order to explore the practicality of our two-step adiabatic quantum approach to finding software errors, we have applied the algorithm to a program of limited size containing a logical error. We did this by calculating the results of the algorithm

assuming perfect adiabatic quantum optimization steps on a processor with few ($N <$ 30) available qubits. Preliminary characterizations of the accuracy achievable using such an algorithm given a set of weak classifiers with certain characteristics are also presented.

6.1 The triplex monitor miscompare problem

The problem we chose to implement is a toy model of program design practices used in mission critical software systems.[5] This program monitors a set of three redundant variables $\{A_t, B_t, C_t\}$ for internal consistency. The variables could represent, e.g., sensor inputs, control signals, or particularly important internal program values. If one value is different from the other two over a predetermined number of snapshots in time t, a problem in the system is indicated and the value of the two consistent redundant variables is propagated as correct. Thus the program is supposed to implement a simple majority-vote error-detection code.

We consider only the simplest case of two time snapshots, i.e., $t = 1, 2$. As just explained, a correct implementation of the monitoring routine should fail a redundant variable A, B, or C if that *same* variable miscompares with both of the other variables in each of the two time frames. The erroneous implemented program we shall consider has the logical error that, due to a mishandled internal implementation of the miscompare tracking over multiple time frames, it fails a redundant variable any time there has been a miscompare in both time frames, even if the miscompare implicated a *different* variable in each time frame.

In order to facilitate quantum V&V using the smallest possible number of qubits, we assume the use of classical preprocessing to reduce the program to its essential structure. The quantum algorithm does not look at the values of the three redundant variables in each time frame. Instead, it sees three logical bits per snapshot, telling it whether each pair of variables is equal. This strategy is also reflected in the program outputs, which are three logical bits indicating whether or not each redundant variable is deemed correct by the monitoring routine. Thus there are nine logical bits, as specified in Table 2.

In terms of Boolean logic, the two behaviors are as follows:
Program Specification

$$x_7 = x_1 \wedge x_3 \wedge x_4 \wedge x_6, \tag{85a}$$
$$x_8 = x_1 \wedge x_2 \wedge x_4 \wedge x_5, \tag{85b}$$
$$x_9 = x_2 \wedge x_3 \wedge x_5 \wedge x_6, \tag{85c}$$

i.e., a variable is flagged as incorrect if and only if it has miscompared with all other variables in all time frames.

---

[5] We are grateful to Greg Tallant from the Lockheed Martin Corporation for providing us with this problem as an example of interest in flight control systems.

Table 2 Logical bits and their significance in terms of variable comparison in the Triplex Miscompare problem

Bit	Significance
x_1	$A_1 \neq B_1$
x_2	$B_1 \neq C_1$
x_3	$A_1 \neq C_1$
x_4	$A_2 \neq B_2$
x_5	$B_2 \neq C_2$
x_6	$A_2 \neq C_2$
x_7	A failed
x_8	B failed
x_9	C failed

Erroneous Program Implementation

$$x_7 = ((x_1 \wedge x_2) \vee (x_2 \wedge x_3) \vee (x_1 \wedge x_3)) \wedge x_4 \wedge x_6, \tag{86a}$$

$$x_8 = ((x_1 \wedge x_2) \vee (x_2 \wedge x_3) \vee (x_1 \wedge x_3)) \wedge x_4 \wedge x_5, \tag{86b}$$

$$x_9 = ((x_1 \wedge x_2) \vee (x_2 \wedge x_3) \vee (x_1 \wedge x_3)) \wedge x_5 \wedge x_6, \tag{86c}$$

i.e., a variable is flagged as incorrect if it miscompares with the other variables in the final time frame and if any variable has miscompared with the others in the previous time frame.

6.2 Implemented algorithm

The challenges before us are to train classifiers to recognize the behavior of both the program specification and the erroneous implementation, and then to use those classifiers to find the errors. These objectives have been programmed into a hybrid quantum-classical algorithm using the quantum techniques described in Sects. 3 and 5 and classical strategy refinements based on characteristics of available resources (for example, the accuracy of the set of available weak classifiers). The performance of this algorithm has been tested through computational studies using a classical optimization routine in place of adiabatic quantum optimization calls.

The algorithm takes as its inputs two training sets, one for the specification classifier and one for the implementation classifier. The two strong classifiers are constructed using the same method, one after the other, consulting the appropriate training set.

When constructing a strong classifier, the algorithm first evaluates the performance of each weak classifier in the dictionary over the training set. Weak classifiers with poor performance, typically those with over 40 % error, are discarded. The resulting, more accurate dictionary is fed piecewise into the quantum optimization algorithm.

Ideally, the adiabatic quantum optimization using the final Hamiltonian (25) would take place over the set of all weak classifiers in the modified, more accurate dictionary. However, the reality of quantum computation for some time to come is that the number of qubits available for processing will be smaller than the number of weak classifiers

in the accurate dictionary. This problem is addressed by selecting random groups of Q classifiers (the number of available qubits) to be optimized together. An initial random group of Q classifiers is selected, the optimal weight vector $\mathbf{q}^{\text{opt}}$ is calculated by classically finding the ground state of H_F, and the weak classifiers which receive weight 0 are discarded. The resulting spaces are filled in with weak classifiers randomly selected from the set of those which have not yet been considered, until all Q classifiers included in the optimization return a weight of 1. This procedure is repeated until all weak classifiers in the accurate dictionary have been considered, at which time the most accurate group of Q generated in this manner is accepted as the strong classifier for the training set in question. Clearly, alternative strategies for combining subsets of Q weak classifiers could be considered, such as genetic algorithms, but this was not attempted here.

Both the specification and implementation strong classifiers are generated in this way, resulting in

$$R_{\mathbf{w}Q}(\mathbf{x}) = \sum_{i=1}^{N} w_i^Q h_i(\mathbf{x}) \tag{87}$$

$$T_{\mathbf{z}Q}(\mathbf{x}) = \sum_{i=1}^{N} z_i^Q h_i(\mathbf{x}) \tag{88}$$

where w_i^Q and z_i^Q take the value 1 if the corresponding weak classifier $h_i(\mathbf{x})$ is selected using the iterative procedure described in the preceding paragraph, and are zero otherwise. This is the same structure as that seen in Eqs. (70) and (71), but with different vectors $\mathbf{w}$ and $\mathbf{z}$ due to the lack of available qubits to perform a global optimization over the accurate dictionary.

The two strong classifiers of Eqs. (87) and (88) are summed as in Eq. (72) to create a final energy function that will push errors to the bottom part of the spectrum. This is translated to a final Hamiltonian H_F as in Eq. (84) and the result of the optimization (i.e., the ground state of this H_F) is returned as the error candidate. This portion of the algorithm makes it crucial to employ intelligent classical preprocessing in order to keep the length of the input and output vectors as small as possible, because each bit in the input–output vector corresponds to a qubit, and the classical cost of finding the ground state of H_F grows exponentially with the number of qubits.

6.3 Simulation results

Our simulation efforts have focused on achieving better accuracy from the two strong classifiers. If the strong classifiers are not highly accurate, the second part of the algorithm, the quantum-parallel use of the classifiers, will not produce useful results because input–output vectors the classifiers do not handle correctly could occupy the low-lying spectrum.

In the interest of pushing the limits of accuracy of the strong classifiers, some simulations were performed on the miscompare problem in a single time frame. Under

Fig. 6 Error fractions in 16-member specification classifier calculations; *Left*: average over 50. *Right*: best of 50


Fig. 7 Error fractions in 16-member implementation classifier calculations; *Left*: average over 50. *Right*: best of 50

this simplification, the program specification and implementation are identical (the error arises over multiple time frames), and indeed the numerical results will show that the results for the two classifiers are the same (see Figs. 6, 7, right).

The algorithm described in Sect. 6.2 was run 50 times, each time producing two strong classifiers comprising 16 or fewer weak classifier members. The figure of 16 qubits was chosen because it allowed the computations to be performed in a reasonable amount of time on a desktop computer while still allowing for some complexity in the makeup of the strong classifiers. This set of 50 complete algorithmic iterations was performed for 26 values of λ, the sparsity parameter introduced in Eq. (18). The average percentage of error for both strong classifiers was examined, as was the best error fraction achieved in the 50 iterations. These two quantities are defined as follows:

$$\mathrm{err}_{\mathrm{avg}} = \frac{1}{50} \sum_{i=1}^{50} L_i(\mathbf{w}^{\mathrm{opt}}) \tag{89}$$

$$\mathrm{err}_{\mathrm{min}} = \min_i L_i(\mathbf{w}^{\mathrm{opt}}), \tag{90}$$

where L is the function that counts the total number of incorrect classifications, Eq. (17). The weight vector $\mathbf{z}^{\mathrm{opt}}$ can be substituted for $\mathbf{w}^{\mathrm{opt}}$ in Eqs. (89) and (90)

Fig. 8 Error fractions of specification (*left*) and implementation (*right*) classifiers, for an increasing number of qubits

if the strong classifier being analyzed is the implementation rather than the specification classifier.

Both the average and minimum error for the specification and implementation classifiers are plotted in Figs. 6 and 7, respectively, as a function of λ.

As shown in Figs. 6 and 7, while the average percent error for both classifiers hovered around 25 %, the best percent error was consistently just below 16 % for both the specification and implementation classifiers. The consistency suggests two things: that the randomness of the algorithm can be tamed by looking for the best outcome over a limited number of iterations, and that the sparsity parameter, λ, did not have much effect on classifier accuracy.

Noting in particular the lack of dependency on λ, we move forward to examine the results of simulations on more difficult and computationally intensive applications of the algorithm. These results address the triplex monitor miscompare problem exactly as described in Sect. 6.1 and increase the number of qubits as far as 26. The error fractions of the best strong classifiers found, defined as

$$\text{err}_{\min}(Q) = \min_i L_i(\mathbf{w}^{\text{opt}}) \qquad i \in \{1, \ldots, n_{\text{sim}}(Q)\} \tag{91}$$

where $n_{\text{sim}}(Q)$ is the number of simulations performed at Q qubits, are plotted in Fig. 8 as a function of the number of qubits allowed in the simulation.

For $Q = 16$ through $Q = 23$, the error fraction shown is for the best-performing classifier, selected from 26 iterations of the algorithm that were calculated using different values of λ. The consistently observed lack of dependence on λ in these and other simulations (such as the 50-iteration result presented above) justifies this choice. For $Q = 24$ to $Q = 26$, it was too computationally intensive to run the algorithm multiple times, even on a high performance computing cluster, so the values plotted are from a single iteration with λ assigned to zero. This was still deemed to be useful data given the uniformity of the rest of the simulation results with respect to λ. The dependence on the parity of the number of qubits is a result of the potential for the strong classifier to return 0 when the number of weak classifiers in the majority vote is even. Zero is not technically a misclassification in that the classifier places the vector

x in the wrong class, but neither does the classifier give the correct class for **x**. Rather, we obtain a " don't-know" answer from the classifier, which we do not group with the misclassifications because it is not an outright error in classification. It is a different, less conclusive piece of information about the proper classification of **x** which may in fact be useful for other applications of such classifiers.

The important conclusion to be drawn from the data quantifying strong classifier errors as a function of the number of available qubits is that performance seems to be improving only slightly as the number of available qubits increases. This may indicate that even with only 16 qubits, if the algorithm is iterated a sufficient number of times to compensate for its random nature, the accuracy achieved is close to the limit of what can be done with the current set of weak classifiers. This is encouraging in the context of strong classifier generation and sets a challenge for improving the performance of weak classifiers or breaking the problem into intermediate stages.

6.4 Comparison of results with theory

In light of the conditions for an ideal strong classifier developed in Sect. 4, it is reasonable to ask the following questions: How close do the weak classifiers we have for the problem studied here come to satisfying the conditions? What sort of accuracy can we expect our simulations to yield? Figure 9 and a few related calculations shed some light on the answers. In the figure, each row of pixels represents a single weak classifier in the dictionary and each column represents one vector in the input–output space. Horizontal red lines divide the different levels of performance exhibited by the weak classifiers. White pixels represent a given weak classifier categorizing a given input–output vector correctly. Black pixels represent incorrect classifications.

The problematic aspect of Fig. 9 is the vertical bars of white and black exhibited by some of the more accurate classifiers. The method detailed above for constructing a completely accurate strong classifier relies on pairs of classifiers which are correct where others fall short, and which do not both classify the same input–output vector incorrectly. This is impossible to find in the most accurate group of weak classifiers alone, given that there are black bars of erroneous classifications spanning the entire height of the set.

For numerical analysis of the performance of the set of Boolean weak classifiers on the sample problem, we relate the statistics of the dictionary on the input–output vector space $\mathcal{V}$ to Conditions 2 and 2a. Three quantities will be useful for this analysis. The first is the error fraction of an individual weak classifier

$$
\eta_j = 1 - \frac{1}{|\mathcal{T}|} \sum_{s=1}^{|\mathcal{T}|} H\left[y_s h_j(\mathbf{x_s})\right], \tag{92}
$$

that is, the fraction of the training set incorrectly classified by the weak classifier $h_j(\mathbf{x})$. We use the Heaviside step function to count the number of vectors correctly classified.

Fig. 9 Accuracy of weak classifier dictionary on input–output vector space. *White/black* pixels represent a weak classifier $h_i(\mathbf{x})$ (all weak classifiers meeting Condition 1 indexed in order of increasing error η_j as in Eq. (92) on *vertical* axis) categorizing an input–output vector (indexed in lexicographical order on *horizontal* axis, there are 2^9 vectors arising from the 9 Boolean variables in the sample problem) correctly/incorrectly, respectively. These classifications were to determine whether an input–output pair was correct or erroneous, i.e., we are analyzing the performance of the specification classifier

Next is the *minimum possible* overlap of correctly classified vectors for a pair of weak classifiers over $\mathcal{V}$:

$$\phi_{jj'} = 1 - \eta_j - \eta_{j'} \qquad (93)$$

In Eq. (93), we add the correctness fraction $(1 - \eta_j)$ of each weak classifier, then subtract 1 to arrive at the number of vectors that must be classified correctly by both weak classifiers at once.

The next definition we shall require is that of the *actual* overlap of correct classifications:

$$\gamma_{jj'} = \frac{1}{|\mathcal{T}|} \sum_{s=1}^{|\mathcal{T}|} H\left[y_s\left(h_j(\mathbf{x_s}) + h_{j'}(\mathbf{x_s})\right)\right] \equiv \phi_{jj'} + \epsilon_{jj'} \tag{94}$$

In Eq. (94), we count the number of vectors that are actually classified correctly by both weak classifiers.

If the minimum possible and actual overlaps are the same, i.e., $\epsilon_{jj'} = 0$, then Condition 2 holds, and the weak classifier pair has minimum *correctness overlap*. Otherwise, if $\phi_{jj'} \neq \gamma_{jj'}$, only the weaker Condition 2a is satisfied, so the weak classifier pair has a greater than minimal correctness overlap and a forced overlap of incorrect classifications $\epsilon_{jj'} > 0$ (see Fig. 2) that could cancel out the correct votes of a different weak classifier pair and cause the strong classifier to be either incorrect or inconclusive.

Our numerical analysis of the weak classifiers satisfying Condition 1 (having $\eta_j <$ 0.5) showed that the average correctness overlap $\gamma_{jj'}$ between any two weak classifiers was 0.3194. The maximum correctness overlap for any pair of weak classifiers was $\gamma_{jj'} = 0.6094$. The minimum was $\gamma_{jj'} = 0.1563$, between two weak classifiers with respective error fractions (amount of the training set misclassified by each individual weak classifier) of $\eta_j = 0.4844$ and $\eta_{j'} = 0.4531$. Compare this to the minimum possible overlap with two such classifiers, $\phi_{jj'} = 0.0625$, and it becomes apparent that this set of weak classifiers falls short of ideal, given that $\epsilon_{jj'} = 0.0938$ for the weak classifier pair with minimum overlap.

When only the most accurate weak classifiers ($\eta_j = 0.3906$; above the top red horizontal line in Fig. 9) were included, the average correctness overlap was $\gamma_{jj'} = 0.4389$, the maximum was $\gamma_{jj'} = 0.6094$, and the minimum was $\gamma_{jj'} = 0.3594$. In order to come up with a generous estimate for the accuracy achievable with this group of weak classifiers, we focus on the minimum observed correctness overlap. The minimum possible correctness overlap for two classifiers with $\eta_j = 0.3906$ is $\phi_{jj'} = 0.2188$. With an ideal set of weak classifiers of error $\eta_j = 0.3906$ and correctness overlap $\phi_{jj'} = 0.2188$, it would take seven weak classifiers to construct a completely accurate strong classifier: three pairs of two classifiers each to cover a fraction 0.6564 of the solution space with a correctness overlap from one of the pairs, and one more weak classifier to provide the extra correct vote on the remaining 0.3436 fraction of the space. Assuming that three pairs of weak classifiers with minimum overlap and optimal relationships to the other weak classifier pairs could be found, there will still be a significant error due to the overlap fractions of the pairs being larger than ideal. In fact, each pair of weak classifiers yields an error contribution of $\epsilon_{jj'} = 0.1406$, guaranteeing that a fraction $3\epsilon_{jj'} = 0.4218$ of the input–output vectors will be classified incorrectly by the resulting strong classifier. This is not far from the simulation results for odd-qubit strong classifiers (Fig. 8, left), which suggests that the algorithm currently in use is producing near-optimal results for the dictionary of weak classifiers it has access to.

7 Conclusions

We have developed a quantum adiabatic machine learning approach and applied it to the problem of training a quantum software error classifier. We have also shown how to use this classifier in quantum-parallel on the space of all possible input–output pairs of a given implemented software program P. The training procedure involves selecting a set of weak classifiers, which are linearly combined, with binary weights, into two strong classifiers.

The first quantum aspect of our approach is an adiabatic quantum algorithm which finds the optimal set of binary weights as the ground state of a certain Hamiltonian. We presented two alternatives for this algorithm. The first, inspired by [6,17], gives weight to single weak classifiers to find an optimal set. The second algorithm for weak classifier selection chooses pairs of weak classifiers to form the optimal set and is based on a set of sufficient conditions for a completely accurate strong classifier that we have developed.

The second quantum aspect of our approach is an explicit procedure for using the optimal strong classifiers in order to search the entire space of input–output pairs in quantum-parallel for the existence of an error in P. Such an error is identified by performing an adiabatic quantum evolution, whose manifold of low-energy final states favors erroneous states.

A possible improvement of our approach involves adding intermediate training spaces, which track intermediate program execution states. This has the potential to fine-tune the weak classifiers, and overcome a limitation imposed by the desire to restrict our Hamiltonians to low-order interactions, yet still account for high-order correlations between bits in the input–output states.

An additional improvement involves finding optimal interpolation paths $s(t)$ (26) from the initial to the final Hamiltonian [52,53], for both the classifier training and classifier implementation problems.

We have applied our quantum adiabatic machine learning approach to a problem with real-world applications in flight control systems, which has facilitated both algorithmic development and characterization of the success of training strong classifiers using a set of weak classifiers involving minimal bit correlations.

References

1. Vapnik, V.N.: Statistical Learning Theory. Wiley, London (1998)
2. Servedio, R.A., Gortler, S.J.: Equivalences and separations between quantum and classical learnability. SIAM J. Comput. **33**, 1067 (2004)
3. Aïmeur, E., Brassard, G., Gambs, S.: Machine learning in a quantum world. In: Lamontagne, L., Marchand, M. (eds.) Advances in Artificial Intelligence, vol. 4013 of Lecture Notes in Computer Science, p. 431. Springer, Berlin (2006)

4. Meir, R., Rätsch, G.: An introduction to boosting and leveraging. In: Mendelson, S., Smola, A. (eds.) Advanced Lectures on Machine Learning, vol. 2600 of Lecture Notes in Computer Science, p. 118. Springer, Berlin (2003)
5. Freund, Y., Schapire, R., Abe, N.: A short introduction to boosting. J. Jpn. Soc. Artif. Intell. **14**, 771 (1999)
6. Neven, H., Denchev, V.S., Rose, G., Macready, W.G.: Training a binary classifier with the quantum adiabatic algorithm. eprint arXiv:0811.0416
7. Neven, H., Denchev., V.S., Drew-Brook, M., Zhang, J., Macready, W.G., Rose, G.: NIPS 2009 demonstration: Binary classification using hardware implementation of quantum annealing (2009)
8. Chandola, V., Banerjee, A., Kumar, V.: Anomaly detection: A survey. ACM Comput. Surv. (CSUR) **41**(3), 15 (2009)
9. Dijkstra, E.W.: Notes on structured programming. In: Dahl, O.-J., Dijkstra, E.W., Hoare, C.A.R. (eds.) Structured Programming, p. 1. Academic Press, New York (1972)
10. Tassey, G.: The economic impacts of inadequate infrastructure for software testing. National Institute of Standards and Technology, RTI Project 7007.011 (2002)
11. Bryce, R., Kuhn, R., Lei, Y., Kacker, R.: Combinatorial testing. In: Ramachandran, M., de Carvalho, R.A. (eds.) Handbook of Software Engineering Research and Productivity Technologies, p. 196. IGI Global (2009)
12. Kuhn, D.R., Kacker, R.N., Lei, Y.: Practical combinatorial testing. NIST Special, Publication 800–142 (2010)
13. Grindal, M., Offutt, J., Andler, S.F.: Combination Testing Strategies: A survey. GMU Technical, Report ISE-TR-04-05 (2004)
14. Cohen, D.M., Dalal, S.R., Parelius, J., Patton, G.C.: The combinatorial design approach to automatic test generation. Softw. IEEE **13**, 83 (1996)
15. D'Silva, V., Kroening, D., Weissenbacher, G.: A survey of automated techniques for formal software verification. IEEE Trans. Comput. Aided Des. Integr. Circuits Syst. **27**, 1165 (2008)
16. Weber, T., Amjad, H.: Efficiently checking propositional refutations in HOL theorem provers. J. Appl. Log. **7**, 26 (2009)
17. Neven, H., Rose, G., Macready, W.G.: Image recognition with an adiabatic quantum computer I. mapping to quadratic unconstrained binary optimization. eprint arXiv:0804.4457
18. Neven, H., Denchev, V.S., Rose, G., Macready, W.G.: Training a large scale classifier with the quantum adiabatic algorithm. eprint arXiv:0912.0779
19. Bian, Z., Chudak, F., Macready, W.G., Rose, G.: The Ising model: teaching an old problem new tricks. D-Wave Systems (2010)
20. Schapire, R.E.: The strength of weak learnability. Mach. Learn. **5**, 197 (1990)
21. Farhi, E., Goldstone, J., Gutmann, S., Sipser, M.: Quantum computation by adiabatic evolution. eprint quant-ph/0001106
22. Farhi, E., Goldstone, J., Gutmann, S., Lapan, J., Lundgren, A., Preda, D.: A quantum adiabatic evolution algorithm applied to random instances of an NP-complete problem. Science **292**(5516), 472 (2001)
23. Aharonov, D., van Dam, W., Kempe, J., Landau, Z., Lloyd, S., Regev, O.: Adiabatic quantum computation is equivalent to standard quantum computation. SIAM J. Comput. **37**, 166 (2007)
24. Mizel, A., Lidar, D.A., Mitchell, M.: Simple proof of equivalence between adiabatic quantum computation and the circuit model. Phys. Rev. Lett. **99**, 070502 (2007)
25. Jordan, S.P., Farhi, E., Shor, P.W.: Error-correcting codes for adiabatic quantum computation. Phys. Rev. A **74**, 052322 (2006)
26. Lidar, Daniel A.: Towards fault tolerant adiabatic quantum computation. Phys. Rev. Lett. **100**, 160506 (2008)
27. Childs, Andrew M., Edward, Farhi, John, Preskill: Robustness of adiabatic quantum computation. Phys. Rev. A **65**, 012322 (2001)
28. Sarandy, M.S., Lidar, D.A.: Adiabatic quantum computation in open systems. Phys. Rev. Lett. **95**, 250503 (2005)
29. Stehle, E., Lynch, K., Shevertalov, M., Rorres, C., Mancoridis, S.: On the use of computational geometry to detect software faults at runtime. ICAC10, June 711. Washington, DC, USA (2010)
30. Le Traon, Y., Baudry, B., Jezequel, J.-M.: Design by contract to improve software vigilance. IEEE Trans. Softw. Eng. **32**, 571 (2006)

31. Mannor, S., Meir, R.: Geometric bounds for generalization in boosting. In: Helmbold, D., Williamson, B. (eds.) Computational Learning Theory, vol. 2111 of Lecture Notes in Computer Science, pp. 461–472. Springer, Berlin (2001)
32. Kotsiantis, S.B.: Supervised machine learning: A review of classification techniques. Informatica **31**, 249 (2007)
33. Yu, L., Liu, H.: Efficient feature selection via analysis of relevance and redundancy. J. Mach. Learn. Res. **5**, 1205 (2004)
34. Zhang, S., Zhang, C., Yang, Q.: Data preparation for data mining. Appl. Artif. Intell. **17**, 375 (2003)
35. Cheng, H., Yan, X., Han, J., Hsu, C.-W.: Discriminative frequent pattern analysis for effective classification. In: International Conference on Data Engineering, p. 716 (2007)
36. Breiman, L.: Arcing classifiers. Ann. Stat. **26**, 801 (1998)
37. Blumer, A., Ehrenfeucht, A., Haussler, D., Warmuth, M.K.: Occam's razor. Inf. Process. Lett. **24**, 377 (1987)
38. Biamonte, J.D., Peter, Love: Realizable Hamiltonians for universal adiabatic quantum computers. Phys. Rev. A **78**, 012352 (2008)
39. Choi, V.: Minor-embedding in adiabatic quantum computation: I. The parameter setting problem. Quantum Inf. Process. **7**, 193 (2008)
40. Karimi, K., Dickson, N.G., Hamze, F., Amin, M.H.S., Drew-Brook, M., Chudak, F.A., Bunyk, P.I., Macready, W.G., Rose, G.: Investigating the performance of an adiabatic quantum optimization processor. Quantum Inf. Process. **11**(1), 77 (2012)
41. Harris, R., Johnson, M.W., Lanting, T., Berkley, A.J., Johansson, J., Bunyk, P., Tolkacheva, E., Ladizinsky, E., Ladizinsky, N., Oh, T., Cioata, F., Perminov, I., Spear, P., Enderud, C., Rich, C., Uchaikin, S., Thom, M.C., Chapple, E.M., Wang, J., Wilson, B., Amin, M.H.S., Dickson, N., Karimi, K., Macready, W., Truncik, C.J.S., Rose, G.: Experimental investigation of an eight-qubit unit cell in a superconducting optimization processor. Phys. Rev. B **82**, 024511 (2010)
42. Cheng, H., Yan, X., Han, J., Hsu, C.-W.: Discriminative frequent pattern analysis for effective classification. In: IEEE 23rd International Conference on Data Engineering, Istanbul, Turkey (2007)
43. Teufel, S.: Adiabatic Perturbation Theory in Quantum Dynamics. Springer, Berlin (2003)
44. Jansen, S., Ruskai, M.-B., Seiler, R.: Bounds for the adiabatic approximation with applications to quantum computation. J. Math. Phys. **48**, 102111 (2007)
45. Lidar, D.A., Rezakhani, A.T., Hamma, A.: Adiabatic approximation with exponential accuracy for many-body systems and quantum computation. J. Math. Phys. **50**, 102106 (2009)
46. Roland, J., Cerf, N.J.: Quantum search by local adiabatic evolution. Phys. Rev. A **65**, 042308 (2002)
47. Rezakhani, A.T., Pimachev, A.K., Lidar, D.A.: Accuracy versus run time in an adiabatic quantum search. Phys. Rev. A **82**, 052305 (2010)
48. Young, A.P., Knysh, S., Smelyanskiy, V.N.: Size dependence of the minimum excitation gap in the quantum adiabatic algorithm. Phys. Rev. Lett. **101**, 170503 (2008)
49. Slepian, D.: On the number of symmetry types of Boolean functions of N variables. Can. J. Math. **5**, 185 (1953)
50. Bryant, R.E.: Graph-based algorithms for Boolean function manipulation. IEEE Trans. Comput. **C–35**, 677 (1986)
51. Jordan, Stephen P., Edward, Farhi: Perturbative gadgets at arbitrary orders. Phys. Rev. A **77**, 062329 (2008)
52. Rezakhani, A.T., Kuo, W.-J., Hamma, A., Lidar, D.A., Zanardi, P.: Quantum adiabatic brachistochrone. Phys. Rev. Lett. **103**, 080502 (2009)
53. Rezakhani, A.T., Abasto, D.F., Lidar, D.A., Zanardi, P.: Intrinsic geometry of quantum adiabatic evolution and quantum phase transitions. Phys. Rev. A **82**, 012321 (2010)