

# QUILT: Effective Multi-Class Classification on Quantum Computers Using an Ensemble of Diverse Quantum Classifiers

Daniel Silver, Tirthak Patel, Devesh Tiwari

Northeastern University  
{silver.da, patel.ti, d.tiwari}@northeastern.edu

## Abstract

Quantum computers can theoretically have significant acceleration over classical computers; but, the near-future era of quantum computing is limited due to small number of qubits that are also error prone. Quilt is a framework for performing multi-class classification task designed to work effectively on current error-prone quantum computers. Quilt is evaluated with real quantum machines as well as with projected noise levels as quantum machines become more noise-free. Quilt demonstrates up to 85% multi-class classification accuracy with the MNIST dataset on a five-qubit system.

## Introduction

**Problem Motivation.** Quantum algorithms have been developed in the areas of chemistry, cryptography, machine learning, and optimization (Lu et al. 2019; Shor 1999; Tiwari and Melucci 2019; Khairy et al. 2020). A class of algorithms known as quantum variational algorithms are designed to optimize and execute quantum machine learning and classification workloads (Benedetti et al. 2019).

While theoretically-promising, existing quantum machine learning classifiers are designed for future large-scale ideal quantum systems. This is because loading data, training, and testing samples on existing near-term intermediate-scale quantum (NISQ) computers is challenging due to significant hardware errors (Schuld and Killoran 2019a; Jurcevic et al. 2021; Preskill 2018). *Consequently, existing quantum classifiers have been demonstrated to be effective only for relatively simple binary classification tasks (Schuld, Fingerhuth, and Petruccione 2017; Grant et al. 2018). As our evaluation confirms, existing state-of-the-art approach are ineffective for multi-class classification (e.g., <30% accuracy for eight-class image classification). Currently, there is lack of capability to perform multi-class classification tasks on the real quantum machines for exploration and improvement.*

**Contributions.** Quilt specifically bridges this gap by open-sourcing its framework and dataset to the community for multi-class classification on NISQ quantum machines. Quilt makes the following key contributions:

(1) A key idea behind Quilt is to build an ensemble of quantum classifiers to perform multi-class classification.

While intuitive and widely-applied in the classical world, we found that classical way of building ensembles is ineffective – it improves the accuracy only by 5%-10% over the state of the art. To mitigate this challenge, Quilt builds a novel ensemble of quantum classifiers by (1) constructing diverse classifiers, and (2) keeping them simple.

(2) Second, we demonstrate that Quilt further enhances the classification accuracy obtained via carefully constructed diverse ensembles of models. Quilt makes an observation that when outputs from diverse models are combined, individual models agree on “most parts” of the output, but they may disagree on “some parts” due to the intrinsic noise, which cannot be mitigated simply by increasing the diversity or number of members. To mitigate this, Quilt designs a safeguard error correcting mechanism that detects which parts have most disagreement and then, builds specific error-correcting binary classifiers to increase the consensus for those parts, improving the overall prediction accuracy.

(3) Third, we show how Quilt performs on a variety of datasets and in different simulation environments. We evaluate Quilt on MNIST, Fashion-MNIST, and Cifar for four and 8 classes. Quilt demonstrates an advantage in every case by up to 46% over state-of-the-art for eight classes of MNIST. We also demonstrate that Quilt can be run on real quantum machines today and test on IBM Manila and IBM Lima.

## Quantum Classification Background

**Qubit and State Manipulation.** The qubit is the fundamental building block of a quantum computer. Where a classical bit is strictly binary, a qubit can be in a *superposition* of both binary states, as shown below in Dirac notation:

$$|\Psi\rangle = \alpha|0\rangle + \beta|1\rangle \text{ s.t. } \alpha, \beta \in \mathbb{C} \ \& \ \|\alpha\|^2 + \|\beta\|^2 = 1 \quad (1)$$

Here,  $|\Psi\rangle$  is the state of the qubit, which is in a superposition of the  $|0\rangle$  and  $|1\rangle$  basis states. When the qubit’s state is measured, its superposition collapses and outputs state  $|0\rangle$  with probability  $\|\alpha\|^2$  and  $|1\rangle$  with probability  $\|\beta\|^2$ .

Qubit states can be manipulated using unitary transformations that act as single- or multi-qubit quantum gates. Single-qubit gates are represented as rotation gates with three angle parameters:  $R_3(\theta_1, \theta_2, \theta_3)$ . Arbitrary superpositions can be achieved by tuning the rotation angles of this  $R_3$  gate. Gates spanning multiple qubits have the ability to *entangle* qubits into a unified quantum system.




Figure 1: State-of-the-art multi-class quantum classifiers have (a) quadratic scaling for number of models and (b) quickly degrading accuracy scaling with number of classes.

**Quantum Circuits.** A quantum circuit is a sequence of gates applied to a system of qubits to affect the outcome of their entangled superposition state. The goal of this paper is to perform classification. On quantum computing systems, a specific type of quantum circuits – quantum variational circuits (QVCs) – have been used for classification tasks in place of classical neural networks due to their more efficient structure using far fewer parameters and their theorized exponential scaling properties over classical algorithms (Chen et al. 2020; Lockwood and Si 2020). Therefore, next we briefly introduce the basic functions of QVCs.

**Quantum Variational Circuits.** QVCs are tunable quantum circuits that can be optimized for a specific purpose through training (Schuld and Killoran 2019b). QVCs are built using a mix of fixed gates (e.g., single-qubit gates with fixed rotation angles), which encode fixed classical data  $x$  (e.g., input features used for a classification task) and tunable gates (e.g., single-qubit gates with parameterized rotation angles), which encode objective function information  $\theta$  (e.g., maximize classification accuracy). A function  $f$  is then created to map these parameters  $(x, \theta)$  to a quantum system. This function can be any series of unitary operations that define the unique circuit architecture, applied to the initial  $|0\rangle$  state.

$$f(x, \theta) = U(x, \theta) |0\rangle \quad (2)$$

The parameters of this function are tuned via a hybrid classical-quantum computing approach for objective optimization (Benedetti et al. 2019; Havlíček et al. 2019). QVCs are fully differentiable and can therefore, be optimized through standard deep learning optimization algorithms.

**QVCs for Classification.** QVCs are a suitable choice for a binary classifier as they provide high accuracy with few parameters, can be optimized with the same procedures as a neural network, and are small enough to fit on quantum machines today. A major challenge however is how to scale these classifiers to multi-class classification. Most real-world classification work done today involves more than two classes, but existing quantum classifiers are focused on solving only binary classification tasks and do not place much emphasis on multi-class classification tasks (Schuld, Fingerhuth, and Petruccione 2017) and (Grant et al. 2018).

## Motivation and Related Works

Quantum computing can have exponential acceleration over classical computers in many problem domains (Shor 1999; Baertschy and Li 2001; Smith et al. 2019; Biamonte et al. 2017; Abrams and Lloyd 1999; Lu et al. 2019). Notably, quantum computing is also expected to prove useful for the field of machine learning and classification (Biamonte et al. 2017; Aaronson 2015; Lockwood and Si 2020; Chen et al. 2020; Yang et al. 2020; Li et al. 2021; Khairy et al. 2020). While new discoveries have led to better quantum machines, a defining limitation of the near-term noisy intermediate-scale quantum (NISQ) computers is the lack of large fault-tolerant machines. One is therefore limited to using only a handful of qubits to perform quantum algorithms. This largely constrains the development of quantum algorithms for real-world applications.

**Related Quantum Classification Efforts** Multiple works have been put forward in the space of quantum image classification. For example, (Gambs 2008), (Monràs, Sentís, and Wittek 2017), and (Wang et al. 2021b) are focused on theoretical algorithms that can only run on large fault-tolerant quantum computers; while promising, these works are not applicable in the current NISQ era as the computers are small and noisy. Other works leverage classical neural networks to handle much of the heavy lifting for quantum classification (Nguyen and Kenyon 2018; Hellstem 2021). Unfortunately, in these cases, the potential for quantum speedup is greatly reduced as the works inherit the limitations of classical networks, such as when dealing with data in quantum entangled states (Grant et al. 2018).

The most relevant and recent efforts in the quantum classification space are (Schuld, Fingerhuth, and Petruccione 2017) and (Grant et al. 2018). These efforts are state-of-the-art effort for classification tasks on NISQ quantum computers. In (Schuld, Fingerhuth, and Petruccione 2017), the authors use a simple distance-based classifier in simulation to classify the Iris dataset by analyzing how it performs for each pair of classes (1v2, 2v3, and 1v3) in separate binary classifications. The more recent work – (Grant et al. 2018) – tackles the same image classification problem of splitting the Iris dataset into tuples, with the addition of analyzing binary splits in the MNIST dataset. They use tensor networks to build their binary classifiers, then evaluate the performance on a real quantum computer. However, both (Schuld, Fingerhuth, and Petruccione 2017) and (Grant et al. 2018) are primarily designed to be effective for binary classification. While they attempt to extend to multi-class classification using binary classifications on every unique two-class pairing in their datasets, this ad-hoc extension approach has several limitations, as we discuss next and our evaluation confirms. Recently, (Wang et al. 2021a) has proposed a framework to address noise in classification circuits.

## Why are the State-of-the-Art Quantum Classifiers Not Effective for Multi-Class Classification?

The state-of-the-art approach for multi-class classification (Grant et al. 2018) essentially relies on splitting a multi-class dataset into many binary datasets and classifying each

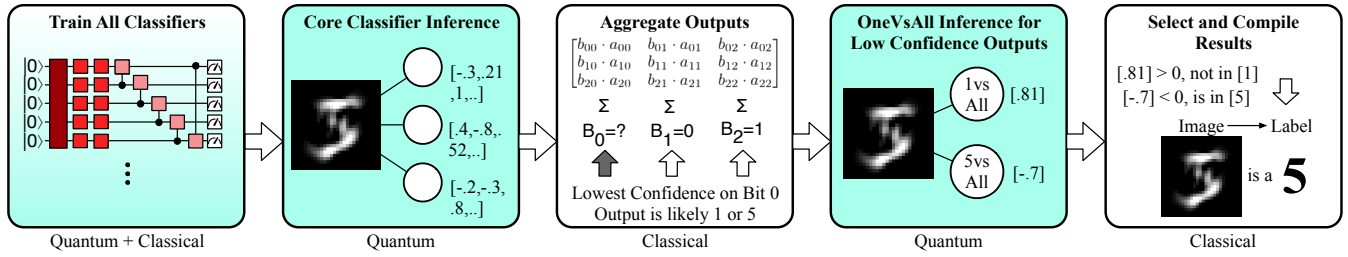


Figure 2: Overview of the design of Quilt. The label under each step indicates if the step is performed quantumly or classically.

pair and combining them all for a multi-class decision – known as the OneVsOne Classifier. For example, for a three-class classification with classes A, B, and C, this would require three unique classifiers; for (A,B), (B,C), and (A,C). Each one votes in favor of one class and the class with two votes wins. However, this method has several shortcomings.

First, *the current approach suffers from poor scalability as a large number of models are needed for training*. The number of models required to build an  $n$ -class multi-class classifier is  $n(n-1)/2$  (Fig. 1(a)). Each model that needs to be trained classically slows down training significantly. Additionally, the inference on quantum machines becomes impractical with larger number of classes. For example, 8 classes would require 28 models, and 16 classes would require 120 different models to be initialized on quantum computers. Instead, *Quilt, as we discuss later, is designed to scale only linearly with the number of classes*.

Second, *the accuracy of the current approach degrades quickly with increase in the number of classes (Fig. 1(b))*. One reason for this is the well-documented issue for the OneVsOne classifier: the “the unclassable region” problem (Galar et al. 2011; Liu, Hao, and Yang 2007). This is where no obvious winner can be selected from the binary classification pairs; for example, a three-class scenario where two candidates tie with one vote each. As more classes are added, the opportunities for ties increase as well. In fact, each additional classifier brings more noise to the system and adds more points of failures if the classifier is biased to a particular class. For this reason, *we design Quilt with an ensemble of classifiers, where each classifier is wholly responsible to classify all classes*.

## Quilt: Design and Implementation

**Overview of the approach.** Fig. 2 illustrates the overview of Quilt. A key idea behind Quilt is to use an ensemble of smaller classifier models instead of using a single large classifier model. This design decision is made to ensure that the impact of quantum operation errors – which are prevalent on quantum real machines – is minimized. Multiple lightweight simple models (referred to as “core classifiers”) make the training and inference process simpler and reduce the impact of errors. In fact, at present, training and inferring large complex models is not feasible due to technological constraints and that is a major roadblock that Quilt curbs.

A small-sized ensemble model is potentially less powerful than a complex, heavyweight model and may produce lower

accuracy results. Therefore, to improve the classification accuracy, Quilt builds an ensemble in a unique way. As shown in the second stage of Fig. 2, each classifier outputs its classification prediction (three bits for an 8-class classification problem) and these inference results are combined across all classifiers by weighting according to their accuracies. As an efficiency measure, the already calculated outputs,  $\hat{x}$ , are used to calculate the accuracy after each iteration. We started with a withheld validation set (10% of the data), but we robustly verified that it did not improve accuracy significantly; so we removed this calculation to save time. As we discuss in detail later, simply replicating the models is not sufficient for improving the accuracy. Quilt specifically designs them to be “different by construction”. Each core classifier can have multiple instances and variants that help diversify the ensemble. Model, variant, and instance design details are discussed later in this section.

During the development of Quilt, we discovered that building a diverse ensemble is also not sufficient. As the fourth stage of Fig. 2 illustrates, when predictions across all classifiers are combined, some bits in the final predicted (e.g., one bit out of three bits for an 8-class prediction task) output may have low-confidence. This is despite careful selection and construction of individual core classifiers. To address this, Quilt designs unique binary classifiers to improve the prediction confidence of low-confidence bits. For example, in Fig. 2, “1 vs All” and “5 vs All” are used because the most significant bit has the lowest confidence and hence, the only output possibility is a ‘1’ or a ‘5’ depending upon the other predicted bits. Finally, Quilt combines the results to predict the target class in a multi-class setting.

Next, before we discuss the model design details of Quilt, we discuss the data processing procedure.

## Preprocessing and Embedding Images and Labels

Running a quantum classifier on a small NISQ quantum machine requires loading and encoding classical data on the quantum computer in a scalable manner. Quilt has selected the model size to be five qubits because models spanning a larger number of qubits increase the surface area of error and hence, risk lowering the overall accuracy.

*Quilt’s solution to encoding each image on a quantum machine is to perform Principal Component Analysis (PCA) on the entire dataset as a preprocessing procedure prior to training the model.* PCA allows particular specification for the exact number of dimensions for data reduction. PCA has




Figure 3: General structure of the different circuit architectures of Quilt.

a one-time offline overhead of processing the image data and generating the 32 features ( $< 1$  minute overhead even with serial processing of 60,000 MNIST images). This allows QUILT to fit on today's NISQ machines and is the most efficient way to encode data to maximize explained variance across a dataset. It is worth noting that other works such as (Henderson et al. 2020) have used threshold based encoding where pixel values above a certain threshold are encoded as the state  $|0\rangle$  and those above are in the state  $|1\rangle$ . While it is possible this may perform well for certain datasets such as MNIST where color and shade are not as important, this should not be expected to perform well for more complex datasets where color is important.

Once the images are preprocessed, Quilt uses amplitude embedding to encode the normalized input vectors on to the qubits. Amplitude embedding transforms  $L$  vectors of normalized classical data of size  $N$  and encodes the data as a quantum state of size  $\log_2(NL)$ . The training set size affects the pre-processing step, but not the characteristics of individual quantum classifiers. The first step of preparing classical data for quantum computation is to initialize all qubits to the ground state:  $|0\rangle$ . Then, the state preparation circuit ( $U_x$ ) is applied to add the fixed classical data to the circuit, embedding one input feature per quantum state amplitude. Once the state is prepared, Quilt's classification model can then be applied. Quilt also preprocesses the labels to ensure compatibility with the quantum space. This requires selecting a value of  $\{-1, 1\}$  for each bit representing the label. With the preprocessing step complete, we delve into the model design details of Quilt.

### Quilt's Model Ensembles

Large and complex models can be used as a building block for classification, but they increase the likelihood of errors on error-prone NISQ machines. To mitigate this challenge, Quilt builds on a base of small-sized circuits because smaller circuit size reduces the adverse impact of errors. In addition, to compensate for the small circuit size, Quilt builds an ensemble of such models that work together to collectively improve the classification performance. Quilt uses accuracy as a weighting mechanism to aggregate the classifications of individual ensemble components. Each ensemble  $\epsilon$  with weight  $w_\epsilon$  contributes to the output  $\hat{y}$  according to the following equation:

$$\hat{y} = \text{sgn} \left[ \sum_{\epsilon} w_{\epsilon} O(x, \theta) \right] \quad (3)$$

where the  $\text{sgn}(z)$  operation is defined as

$$\text{sgn}(z) = \begin{cases} -1, & \text{for } z < 0 \\ 1, & \text{for } z > 0 \end{cases}$$

and the output  $O(x, \theta)$  is a real number between -1 and 1, with its absolute value suggesting the confidence of being a member of the -1 or 1 class. The output is the result of applying the  $U(x, \theta)$  operator to the  $|0\rangle$  state and measuring it, where  $U$  is the unitary transformation applied to input  $x$  and parameters  $\theta$  (from Eq. 2).

The ensemble representation for a QVC then becomes

$$\hat{y} = \text{sgn} \left[ \sum_{\epsilon} \frac{a_{\epsilon}}{a} O(x, \theta) \right] \quad (4)$$

where  $a_{\epsilon}$  is the associated accuracy of ensemble member  $\epsilon$  and  $a$  is the sum of all ensemble accuracies.

### Model Construction

**Core Classifier.** The building block of Quilt's ensemble is the core classifier, which is designed to be able to classify all classes on its own (as opposed an ensemble of OneVsOne classifiers used for state-of-the-art multi-class classification). It consists of an amplitude embedding circuit, followed by repeating rotation gates on each qubit, then CNOT gates to entangle all qubits together, followed by a measurement on all output qubits.

In classical computing, traditional neural networks perform classification by assigning a single class to an output neuron to then use a function such as Softmax to find the most likely candidate. Quilt uses a different approach by leveraging the quantum nature of the model: it uses each measurement qubit  $b$  as a bit in the full  $n$ -bit classification. This allows better scaling to handle a larger number of classes. All measurement qubits form a single bitstring that has an associated class label:

$$\hat{y} = \sum_{b=0}^n 2^{n-b} \left( \frac{\text{sgn}(q_b) + 1}{2} \right) \quad (5)$$

Here,  $q_b$  is the value of the  $b^{\text{th}}$  bit as measured in the quantum state output (Eq. 4). The number of these models scale logarithmically with the number of classes in the label space  $n = \log_2 N$ , but the accuracy scales inversely. This is partially because in order for the entire output to be correct, all the bits have to be correct. For example, if each bit is correct with probability .95, the probability of a correct guess for a three-bit classifier would be  $.95^3 = .857$  or 85.7%.



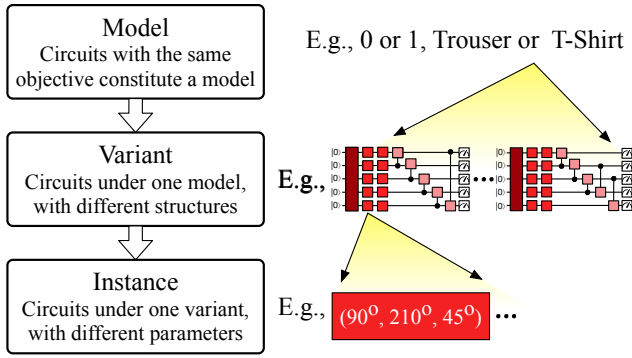


Figure 4: Ensemble diversity hierarchy. At the top of the hierarchy is the "model", which represents a quantum circuit with a unique objective. An example of this would be a OneVsOne classifier with the objective of distinguishing Trousers and T-Shirts. A model can have multiple "variants", which are circuits with the same objective, but different structures. Finally, a variant can have multiple "instances". All instances associated with a variant have the same structure, but may differ in their parameters due to different initial conditions used for optimization.

**Multiple Diverse Variants of Core Classifiers.** Now that we have the general structure for a core classifier defined, our next task is to use this general structure to form an ensemble. It is imperative to diversify the ensemble components as much as possible in order to ensure that they correct for each other's deficiencies, blind spots, and biases. Initializing model instances with different initial conditions enforces some degree of diversity. However, these instances would still be optimizing over the same search space due to the model structure being the same, thus resulting in similar outputs and reduced ensemble diversity.

Quilt mitigates this challenge by constructing model variants with heterogeneous architectures in order to establish more diversity and independence between them. These variants differ in that each variant under the same model can have different number of additional *CNOT* gates and corresponding connections – Fig. 3 provides examples of different variants of the core classifier circuit. Each heterogeneous variant therefore has a different hyper-dimensional search space to optimize over and a different set of minima, resulting in a diverse set of core classifiers with different structures and parameters. The output of these variants can be combined by integrating Quilt's bitwise core classifiers with the weighted ensemble averaging, as expressed below:

$$\hat{y} = \sum_{b=0}^n 2^{n-b} \left[ \frac{\text{sgn}\left(\sum_{\epsilon} \frac{a_{\epsilon}}{a} q_{b\epsilon}\right) + 1}{2} \right] \quad (6)$$

Fig. 4 provides a visual representation of the diversity hierarchy of a model, a variant, and an instance. Different models are the most diverse due to them optimizing different objectives. However, for Quilt's ensemble, we need all core classifier models to have the same objective. Therefore, Quilt employs different variants that are part of the same model

(i.e., have the same objective) and provide better diversity than just having different instances.

**OneVsAll Classifiers for Error Correcting.** While core classifiers do most of the heavy lifting with classification, the addition of support networks improves the overall accuracy of the classification. *One vs all (OneVsAll) classifiers are used for each class in order to provide a parity bit when necessary. These act as error correcting procedures that activate under certain conditions.* A OneVsAll classifier is a binary classifier where one class is assigned to one class and every other class is grouped together in the other class. In the case of Quilt, there is one OneVsAll classifier for each class in the class space. These classifiers only act when the core classifiers have low confidence, defined by a confidence threshold,  $\gamma$ . This is shown in Eq. 7 and is applied to each image separately.

$$\left| \sum_{\epsilon} \frac{a_{\epsilon}}{a} O(x, \theta) \right| < \gamma \quad (7)$$

We find the value of  $\gamma$  is optimal when it is met for 10% of all guesses. For all images, if any have an absolute sum of weighted outputs that falls within  $\gamma$ , we identify the least confident bit  $b$  in the output as:

$$b = \underset{b}{\operatorname{argmin}} \left| \sum_{\epsilon} \frac{a_{\epsilon}}{a} O(x, \theta) \right| \quad (8)$$

The next step is to evaluate the associated OneVsAll classifier for a final verdict on the bit in question. As there is a single bit in question, two OneVsAll Classifiers that are charged with the decision. For example, in the case of evaluating the most significant bit where the two least significant bits are certain (e.g., X01), the output is either a 1 or a 5, so the 1VsAll and 5VsAll classifiers must be evaluated. The bit is only set if and only if both OneVsAll Classifiers being evaluated agree on a decision, otherwise it defaults to its original state. This is shown for classifiers  $C_1$  and  $C_2$  below.

$$q_b = \begin{cases} 0, & \text{for } C_1 = -1, C_2 = 1 \\ 1, & \text{for } C_1 = 1, C_2 = -1 \\ q_b, & \text{otherwise} \end{cases}$$

Here,  $-1$  indicates that the image is in the specific OneVsAll class, and  $1$  indicates that it is not in the class.

### Stitching Together the Ensemble

Overall, Quilt requires five core classifiers +  $n$  OneVsAll classifiers to build an ensemble for  $n$  classes. For example, for its eight-class ensemble, Quilt requires 13 classifiers: five core classifiers (each is a different variant) and eight OneVsAll classifiers. *The number of models required by Quilt therefore scales linearly with the number of classes, as opposed to quadratically with the state-of-the-art approach (Fig. 1).* Quilt's models construct a robust scalable ensemble with error correction built in. Next, we describe the implementation of our training and inference procedure.

### Training and Inference Procedure

Quilt begins by constructing the core model variants, each set to a learning rate of .05, optimized using the Adam optimizer (Kingma and Ba 2014). This is followed by randomly

initializing all variants’ weights separately between 0 and 1. For 100 epochs, these variants sample 50 random images and corresponding labels from the training set split. All vectorized images are passed into the network for evaluation. Within the context of the hybrid quantum-classical framework, gradients are computed classically after generating the outputs of the circuit. The new gate parameters are then optimized by shifting along the gradient to estimate the eigenvalues, then repeating until the optimization has terminated. Quilt uses a form of a squared  $L_2$  loss to perform gradient descent optimization. We find that squared loss works the best in training these small classification networks and outperforms  $L_1$  loss or even techniques enforcing sparsity, such as the  $L_0$  norm. Quilt incorporates  $L_2$  loss in multiclass classification using  $N$  outputs qubits as follows:

$$\text{Loss} = \sum_{i=0}^{\text{batch}} \sum_{j=0}^N (\text{label}[i][j] - \text{prediction}[i][j])^2 \quad (9)$$

As an instance, in the eight-class case, this translates to only training the optimizer against the outputs of the first three qubits and ignoring the last two. As training these types of sensitive quantum networks does not provide a reliable advantage from one epoch to the next, we save the weights for the running best model.

To train the OneVsAll classifiers, we sample the labels to create an equal amount of data in both classes of the classifier. For example, to classify a “7” with 1000 images, the training set would consist of 1000 images of “7” and 1000 images composed of all other digits. These classifiers are also optimized using the Adam optimizer with  $L_2$  loss.

In terms of inference, Quilt runs the five trained core classifiers for each inference task and takes the weighted average of their output. If it determines that there exist bits with low confidence, Quilt runs the respective OneVsAll classifiers for the bit with the lowest confidence. Once the final output is produced, Quilt assesses all the classifier outputs and generates the final label for the inference task.

## Evaluation

### Experimental Methodology

**Dataset.** We evaluate Quilt with the MNIST (Deng 2012), Fashion-MNIST (Xiao, Rasul, and Vollgraf 2017), and Cifar (Krizhevsky, Hinton et al. 2009) datasets, using an 80/20 training-to-testing split. MNIST is a greyscale digit dataset that is widely used for quantum classification (Grant et al. 2018; Wilson et al. 2018; Kerenidis and Luongo 2020; Li, Song, and Wang 2021; Tiwari and Melucci 2019). We also use Fashion-MNIST and Cifar due to the greater difficulty of classifying them than MNIST. For example, Cifar has more pixels than MNIST and is also in RGB. We use classes  $\{(0, 1), 6, 7\}$ ,  $\{2, 3, 4, 5\}$  for MNIST,  $\{(t\text{-shirt/top, trouser, pullover, dress}), \text{coat, sandal, shirt, sneaker}\}$  for Fashion-MNIST, and  $\{(airplane, automobile), \text{bird, cat}\}$ , deer, dog, frog, horse for Cifar. Here, the class sets used for two-class, four-class, and eight-class classification are enclosed in  $()$ ,  $\{\}$ , and  $[\]$ , respectively. We refer to these classes using the naming convention “<dataset>-<classes>” (e.g., Cifar-8 for eight-class Cifar).




Figure 5: Quilt scales better to higher number of classes than comparative techniques for the MNIST dataset.

**Setup Details.** Our solution environment uses Python3, with PennyLane (Bergholm et al. 2018) as a framework using Qiskit (Abraham et al. 2019) as the backbone. We use Sklearn to perform PCA preprocessing for our datasets. Our noise-free simulations take place using the default PennyLane backended simulator. For our noisy simulations, we use “qiskit.aer” to set the error rates of quantum gates. Our real machine evaluations are run on IBM’s five-qubit Lima and Manila computers.

**Comparative Techniques.** (1) **OneVsOne:** As introduced in Fig. 1, this current state-of-the-art technique performs multi-class classification by combining many binary classifiers for a multi-class decision. This is the extension of (Grant et al. 2018). (2) **Ensemble:** The ensemble is similar to Quilt, but without the OneVsAll classifiers or multiple variants. Each member is responsible for a bit of the classification, allowing the binary classifier to scale to multi-class. There are five of these members for each bit and the decision procedure is a simple majority rules. *We include this method for a comparison to show why off-the-shelf ensemble methods are not sufficient for multi-class classification.* We use the same PCA-based dataset preprocessing technique as Quilt for the competitive techniques for fair comparison.

## Results and Analysis

**Quilt achieves higher accuracy than comparative state-of-the-art techniques as the number of classes increases.** We show classification accuracy for MNIST-2, MNIST-4, and MNIST-8 in Fig. 5. We make several observations. We note that all methods perform similarly on 2 classes as binary classification is a relatively simpler task. However, Quilt provides 22% and 46% point better classification accuracy over the state-of-the-art OneVsOne method as the number of classes increases to four and eight. This is because the accuracy-weighted ensemble allows Quilt to factor in model accuracy, allowing more accurate models to be favored. Additionally, Quilt consolidates classification to an ensemble, where each model is responsible for classifying all classes on its own, therefore making better use of the full capacity of the QVC: *where Quilt only requires 13 circuits for eight-class classification, the state-of-the-art OneVsOne classifier requires 28.* However, having an ensemble is not enough; by checking the error-prone bits, Quilt is able to effectively treat one of the bits in each bad sample and obtain higher accuracy than the ensemble method. Even though the simple ensemble approach requires fewer models (i.e., 5), we




Figure 6: Quilt’s accuracy is better than other methods for all well-known (a) four-class and (b) eight-class datasets.




Figure 7: An ensemble with different variants outperforms one with multiple instances of one variant (MNIST-4).

find Quilt’s linear increase in the number of models with the number of classes to be justified given the significant increase in accuracy.

**Quilt outperforms competitive techniques across multiple well-known deep learning image datasets.** In Fig. 6(a), we evaluate how well Quilt performs on Fashion-MNIST-4 and Cifar-4, in addition to MNIST-4. We also look at how Quilt compares to other solutions. Quilt outperforms the state-of-the-art by 46%, 38%, and 14% for MNIST-8, Fashion-MNIST-8, and Cifar-8, respectively. All methods perform better on Fashion-MNIST-4 and MNIST-4 than on Cifar-4. This is because of the limited size of the encoded data that makes encoding the RGB images in Cifar-4 into 32 features tougher (i.e., more information is lost by performing PCA on Cifar images). We also evaluate performance of different datasets for eight classes in Fig. 6(b) and observe similar behavior. Comparing the two figures demonstrates that Quilt is the most scalable and highest performing solution across different datasets.

**Quilt’s strategy of using multiple variants achieves significantly higher accuracy than using a single variant.** Fig. 7 confirms Quilt’s critical design component: higher accuracy is achieved by adding diversity to variant architectures in ensembles. While different instances of a homogeneous variant can be diversified using different initial weights, because these instances traverse the same hyper-dimensional space during parameter optimization, they cannot provide the level of diversity of heterogeneous variants.

**Quilt outperforms other techniques on running image classification tasks on today’s noisy real quantum machines.** Fig. 8 (a) and (b) show Quilt’s classification accu-




Figure 8: (a) MNIST-4 and (b) MNIST-8 accuracy on IBM’s quantum machines Lima and Manila along an interpolation based on half of Manila’s noise levels. OneVsOne could not be computed for eight classes on the real machines as it required an impractical number of model evaluations.




Figure 9: Projected increase in MNIST-8 accuracy as the hardware noise decreases with evolution in hardware.

racy for MNIST-4 and MNIST-8 datasets on real quantum computers. Quilt outperforms other methods due to the redundancies built into it that were designed to make it work well on small and noisy circuits. While we do not focus our optimizations on the noisy machines of today, we still abide by the constraints that allow Quilt to be executable on today’s quantum machines. We also note that the overall accuracy on real systems is lower than our previous simulated results due to the hardware noise. However, in the future, the noise is expected to reduce significantly, with estimates such as  $10^{-4}$  error rate (Jurcevic et al. 2021).

**Quilt’s accuracy performance will only improve with reduction in hardware noise levels as quantum technology advances.** We simulate several noise levels between the ideal case of 0 and the current case of real machines and show how different levels of noise affect the accuracy. Fig. 9 demonstrates that as noise levels continue to improve on real NISQ devices, the accuracy of Quilt improves.

## Conclusion

In this paper, we proposed Quilt, an ensemble-based end-to-end multi-class image classifier for NISQ machines. We showed the scalability potential of Quilt compared to other methods and demonstrated a significant improvement over the state-of-the-art in terms of accuracy and resource requirements. We also showed that these models can be run on today’s quantum machines and project the improvements of QUILT as quantum machines become more error-resistant.

## Acknowledgements

We are thankful for the thoughtful feedback from our anonymous reviewers and the support from Northeastern University, NSF Award 1910601, and the Massachusetts Green High Performance Computing Center (MGHPCC) facility. We acknowledge the use of the IBM Q for this work. The views expressed are those of the authors and do not reflect the official policy or position of IBM or the IBM Q team.

## References

- Aaronson, S. 2015. Read the fine print. *Nature Physics*, 11(4): 291–293.
- Abraham, H.; AduOffei; Agarwal, R.; Akhalwaya, I. Y.; Aleksandrowicz, G.; Alexander, T.; Amy, M.; Arbel, E.; Arijit02; Asfaw, A.; Avkhadiiev, A.; Azaustre, C.; AzizNgoueya; Banerjee, A.; Bansal, A.; Barkoutsos, P.; Barnawal, A.; Barron, G.; Barron, G. S.; Bello, L.; Ben-Haim, Y.; Bevenius, D.; Bhobe, A.; Bishop, L. S.; Blank, C.; Bolos, S.; Bosch, S.; Brandon; Bravyi, S.; Bryce-Fuller; Bucher, D.; Burov, A.; Cabrera, F.; Calpin, P.; Capelluto, L.; Carballo, J.; Carrascal, G.; Chen, A.; Chen, C.-F.; Chen, E.; Chen, J. C.; Chen, R.; Chow, J. M.; Churchill, S.; Claus, C.; Clauss, C.; Cocking, R.; Correa, F.; Cross, A. J.; Cross, A. W.; Cross, S.; Cruz-Benito, J.; Culver, C.; Córcoles-Gonzales, A. D.; Dague, S.; Dandachi, T. E.; Daniels, M.; Dartailh, M.; DavideFrr; Davila, A. R.; Dekusar, A.; Ding, D.; Doi, J.; Drechsler, E.; Drew; Dumitrescu, E.; Dumon, K.; Duran, I.; EL-Safty, K.; Eastman, E.; Eberle, G.; Eendebak, P.; Egger, D.; Everitt, M.; Fernández, P. M.; Ferrera, A. H.; Fouilland, R.; FranckChevallier; Frisch, A.; Fuhrer, A.; Fuller, B.; GEORGE, M.; Gacon, J.; Gago, B. G.; Gambella, C.; Gambetta, J. M.; Gammanpila, A.; Garcia, L.; Garg, T.; Garion, S.; Gilliam, A.; Giridharan, A.; Gomez-Mosquera, J.; Gonzalo; de la Puente González, S.; Gorzinski, J.; Gould, I.; Greenberg, D.; Grinko, D.; Guan, W.; Gunnels, J. A.; Haglund, M.; Haide, I.; Hamamura, I.; Hamido, O. C.; Harkins, F.; Havlicek, V.; Hellmers, J.; Herok, Ł.; Hillmich, S.; Horii, H.; Howington, C.; Hu, S.; Hu, W.; Huang, J.; Huisman, R.; Imai, H.; Imamichi, T.; Ishizaki, K.; Iten, R.; Itoko, T.; JamesSeaward; Javadi, A.; Javadi-Abhari, A.; Javed, W.; Jessica; Jivrajani, M.; Johns, K.; Johnston, S.; Jonathan-Shoemaker; K, V.; Kachmann, T.; Kale, A.; Kanazawa, N.; Kang-Bae; Karazeev, A.; Kassebaum, P.; Kelso, J.; King, S.; Knabberjoe; Kobayashi, Y.; Kovyrshin, A.; Krishnakumar, R.; Krishnan, V.; Krsulich, K.; Kumkar, P.; Kus, G.; LaRose, R.; Lacal, E.; Lambert, R.; Lapeyre, J.; Latone, J.; Lawrence, S.; Lee, C.; Li, G.; Liu, D.; Liu, P.; Maeng, Y.; Majmudar, K.; Malyshev, A.; Manela, J.; Marecek, J.; Marques, M.; Maslov, D.; Mathews, D.; Matsuo, A.; McClure, D. T.; McGarry, C.; McKay, D.; McPherson, D.; Meesala, S.; Metcalfe, T.; Mevissen, M.; Meyer, A.; Mezzacapo, A.; Midha, R.; Minev, Z.; Mitchell, A.; Moll, N.; Montanez, J.; Monteiro, G.; Mooring, M. D.; Morales, R.; Moran, N.; Motta, M.; MrF; Murali, P.; Muggenburg, J.; Nadlinger, D.; Nakanishi, K.; Nannicini, G.; Nation, P.; Navarro, E.; Naveh, Y.; Neagle, S. W.; Neuweiler, P.; Nicander, J.; Niroula, P.; Norlen, H.; NuoWenLei; O’Riordan, L. J.; Ogunbayo, O.; Ollitrault, P.; Otaolea, R.; Oud, S.; Padilha, D.; Paik, H.; Pal, S.; Pang, Y.; Pascuzzi, V. R.; Perriello, S.; Phan, A.; Piro, F.; Pistoia, M.; Piveteau, C.; Pocreau, P.; Pozas-iKerstjens, A.; Prokop, M.; Prutyantov, V.; Puzzuoli, D.; Pérez, J.; Quintiii; Rahman, R. I.; Raja, A.; Ramagiri, N.; Rao, A.; Raymond, R.; Redondo, R. M.-C.; Reuter, M.; Rice, J.; Riedemann, M.; Rocca, M. L.; Rodríguez, D. M.; RohithKarur; Rossmannek, M.; Ryu, M.; SAPV, T.; SamFerracin; Sandberg, M.; Sandesara, H.; Sapra, R.; Sargsyan, H.; Sarkar, A.; Sathaye, N.; Schmitt, B.; Schnabel, C.; Schoenfeld, Z.; Scholten, T. L.; Schoute, E.; Schwarm, J.; Sertage, I. F.; Setia, K.; Shammah, N.; Shi, Y.; Silva, A.; Simonetto, A.; Singstock, N.; Siraichi, Y.; Sitdikov, I.; Sivarajah, S.; Sletfjerding, M. B.; Smolin, J. A.; Soeken, M.; Sokolov, I. O.; Sokolov, I.; SooluThomas; Starfish; Steenken, D.; Stypulkoski, M.; Sun, S.; Sung, K. J.; Takahashi, H.; Takawale, T.; Tavernelli, I.; Taylor, C.; Teylour, P.; Thomas, S.; Tillet, M.; Tod, M.; Tomasik, M.; de la Torre, E.; Trabing, K.; Treinish, M.; TrishaPe; Tuls, D.; Turner, W.; Vaknin, Y.; Valcarce, C. R.; Varchon, F.; Vazquez, A. C.; Villar, V.; Vogt-Lee, D.; Vuillot, C.; Weaver, J.; Weidenfeller, J.; Wiecek, R.; Wildstrom, J. A.; Winston, E.; Woehr, J. J.; Woerner, S.; Woo, R.; Wood, C. J.; Wood, R.; Wood, S.; Wood, S.; Wootton, J.; Yeralin, D.; Yonge-Mallo, D.; Young, R.; Yu, J.; Zachow, C.; Zdan-ski, L.; Zhang, H.; Zoufal, C.; Zoufal; a kapila; a matsuo; beammorison; brandhsn; nick bronn; brosand; chlorophyll zz; csseifms; dekel.meirom; dekelmeirom; dekol; dime10; drholm; dtrenev; ehchen; elfrocampeador; faisaldebouni; fanizzamarco; gabrieleagl; gadial; galeinston; georgios ts; gruu; hhorii; hykavitha; jagunther; jliu45; jscott2; kanejess; klinvill; krutik2966; kurarr; lerongil; ma5x; merav aha-roni; michelle4654; ordmoj; sagar pahwa; rmoyard; saswati qiskit; scottkelso; sethmerkel; shaashwat; sternparky; strick-roman; sumitpuri; tigerjack; toural; tsura crisaldo; vvil-pas; welien; willhbang; yang.luh; yotamvakninibm; and Čepulkovskis, M. 2019. Qiskit: An Open-source Framework for Quantum Computing.
- Abrams, D. S.; and Lloyd, S. 1999. Quantum Algorithm Providing Exponential Speed Increase for Finding Eigenvalues and Eigenvectors. *Phys. Rev. Lett.*, 83: 5162–5165.
- Baertschy, M.; and Li, X. 2001. Solution of a Three-Body Problem in Quantum Mechanics Using Sparse Linear Algebra on Parallel Computers. In *SC ’01: Proceedings of the 2001 ACM/IEEE Conference on Supercomputing*, 31–31.
- Benedetti, M.; Lloyd, E.; Sack, S.; and Fiorentini, M. 2019. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4): 043001.
- Bergholm, V.; Izaac, J.; Schuld, M.; Gogolin, C.; Alam, M. S.; Ahmed, S.; Arrazola, J. M.; Blank, C.; Delgado, A.; Jahangiri, S.; et al. 2018. PennyLane: Automatic differentiation of hybrid quantum-classical computations. *arXiv preprint arXiv:1811.04968*.
- Biamonte, J.; Wittek, P.; Pancotti, N.; Rebentrost, P.; Wiebe, N.; and Lloyd, S. 2017. Quantum machine learning. *Nature*, 549(7671): 195–202.
- Chen, S. Y.-C.; Yang, C.-H. H.; Qi, J.; Chen, P.-Y.; Ma, X.; and Goan, H.-S. 2020. Variational quantum circuits for deep reinforcement learning. *IEEE Access*, 8: 141007–141024.



- Deng, L. 2012. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6): 141–142.
- Galar, M.; Fernández, A.; Barrenechea, E.; Bustince, H.; and Herrera, F. 2011. An overview of ensemble methods for binary classifiers in multi-class problems: Experimental study on one-vs-one and one-vs-all schemes. *Pattern Recognition*, 44(8): 1761–1776.
- Gambs, S. 2008. Quantum classification. *arXiv:0809.0444*.
- Grant, E.; Benedetti, M.; Cao, S.; Hallam, A.; Lockhart, J.; Stojevic, V.; Green, A. G.; and Severini, S. 2018. Hierarchical quantum classifiers. *npj Quantum Information*, 4(1): 1–8.
- Havlíček, V.; Córcoles, A. D.; Temme, K.; Harrow, A. W.; Kandala, A.; Chow, J. M.; and Gambetta, J. M. 2019. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747): 209–212.
- Hellstem, G. 2021. Hybrid Quantum Network for classification of finance and MNIST data. In *2021 IEEE 18th International Conference on Software Architecture Companion (ICSAC-C)*, 1–4.
- Henderson, M.; Shakya, S.; Pradhan, S.; and Cook, T. 2020. Quantvolutional neural networks: powering image recognition with quantum circuits. *Quantum Machine Intelligence*, 2(1): 1–9.
- Jurcevic, P.; Javadi-Abhari, A.; Bishop, L. S.; Lauer, I.; Bogorin, D. F.; Brink, M.; Capelluto, L.; Günlük, O.; Itoko, T.; Kanazawa, N.; Kandala, A.; Keefe, G. A.; Krsulich, K.; Landers, W.; Lewandowski, E. P.; McClure, D. T.; Nannicini, G.; Narasgond, A.; Nayfeh, H. M.; Pritchett, E.; Rothwell, M. B.; Srinivasan, S.; Sundaresan, N.; Wang, C.; Wei, K. X.; Wood, C. J.; Yau, J.-B.; Zhang, E. J.; Dial, O. E.; Chow, J. M.; and Gambetta, J. M. 2021. Demonstration of quantum volume 64 on a superconducting quantum computing system. *Quantum Science and Technology*, 6(2): 025020.
- Kerenidis, I.; and Luongo, A. 2020. Classification of the MNIST data set with quantum slow feature analysis. *Phys. Rev. A*, 101: 062327.
- Khairy, S.; Shaydulin, R.; Cincio, L.; Alexeev, Y.; and Balaprakash, P. 2020. Learning to Optimize Variational Quantum Circuits to Solve Combinatorial Problems. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(03): 2367–2375.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Li, G.; Song, Z.; and Wang, X. 2021. VSQ: Variational Shadow Quantum Learning for Classification. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(9): 8357–8365.
- Li, Q.; Gkoumas, D.; Sordoni, A.; Nie, J.-Y.; and Melucci, M. 2021. Quantum-inspired Neural Network for Conversational Emotion Recognition. *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(15): 13270–13278.
- Liu, B.; Hao, Z.; and Yang, X. 2007. Nesting algorithm for multi-classification problems. *Soft Computing*, 11(4): 383–389.
- Lockwood, O.; and Si, M. 2020. Reinforcement Learning with Quantum Variational Circuit. *Proceedings of the AAAI Conference on Artificial Intelligence and Interactive Digital Entertainment*, 16(1): 245–251.
- Lu, C.; Liu, Q.; Wang, C.; Huang, Z.; Lin, P.; and He, L. 2019. Molecular Property Prediction: A Multilevel Quantum Interactions Modeling Perspective. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 1052–1060.
- Monrás, A.; Sentís, G.; and Wittek, P. 2017. Inductive Supervised Quantum Learning. *Phys. Rev. Lett.*, 118: 190503.
- Nguyen, N. T.; and Kenyon, G. T. 2018. Image Classification Using Quantum Inference on the D-Wave 2X. In *2018 IEEE International Conference on Rebooting Computing (ICRC)*, 1–7.
- Preskill, J. 2018. Quantum Computing in the NISQ era and beyond. *Quantum*, 2: 79.
- Schuld, M.; Fingerhuth, M.; and Petruccione, F. 2017. Implementing a distance-based classifier with a quantum interference circuit. *EPL (Europhysics Letters)*, 119(6): 60002.
- Schuld, M.; and Killoran, N. 2019a. Quantum machine learning in feature Hilbert spaces. *Physical review letters*, 122(4): 040504.
- Schuld, M.; and Killoran, N. 2019b. Quantum Machine Learning in Feature Hilbert Spaces. *Physical Review Letters*, 122(4).
- Shor, P. W. 1999. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM review*, 41(2): 303–332.
- Smith, A.; Kim, M. S.; Pollmann, F.; and Knolle, J. 2019. Simulating quantum many-body dynamics on a current digital quantum computer. *npj Quantum Information*, 5(1).
- Tiwari, P.; and Melucci, M. 2019. Binary Classifier Inspired by Quantum Theory. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33(01): 10051–10052.
- Wang, H.; Ding, Y.; Gu, J.; Lin, Y.; Pan, D. Z.; Chong, F. T.; and Han, S. 2021a. Quantumnas: Noise-adaptive search for robust quantum circuits. *arXiv preprint arXiv:2107.10845*.
- Wang, X.; Ma, Y.; Hsieh, M.-H.; and Yung, M.-H. 2021b. Quantum speedup in adaptive boosting of binary classification. *Science China Physics, Mechanics & Astronomy*, 64(2): 1–10.
- Wilson, C.; Otterbach, J.; Tezak, N.; Smith, R.; Polloreno, A.; Karalekas, P. J.; Heidel, S.; Alam, M. S.; Crooks, G.; and da Silva, M. 2018. Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers. *arXiv preprint arXiv:1806.08321*.
- Xiao, H.; Rasul, K.; and Vollgraf, R. 2017. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*.
- Yang, F.; Jiang, J.; Zhang, J.; and Sun, X. 2020. Revisiting Online Quantum State Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04): 6607–6614.