

IMAGE CLASSIFICATION BY THROWING QUANTUM KITCHEN SINKS AT TENSOR NETWORKS

Anonymous authors

Paper under double-blind review

ABSTRACT

Several variational quantum circuit approaches to machine learning have been proposed in recent years, with one promising class of variational algorithms involving tensor networks operating on states resulting from *local* feature maps. In contrast, a random feature approach known as *quantum kitchen sinks* provides comparable performance, but leverages *non-local* feature maps. Here we combine these two approaches by proposing a new circuit ansatz where a tree tensor network coherently processes the non-local feature maps of quantum kitchen sinks, and we run numerical experiments to empirically evaluate the performance of the new ansatz on image classification. From the perspective of classification performance, we find that simply combining quantum kitchen sinks with tensor networks yields no qualitative improvements. However, the addition of *feature optimization* greatly boosts performance, leading to state-of-the-art quantum circuits for image classification, requiring only shallow circuits and a small number of qubits – both well within reach of near-term quantum devices.

1 INTRODUCTION

Tensor network (TN) methods have been studied across physics, mathematics, and computer science for their expressive power, interpretability, and computational efficiency (Kolda & Bader, 2009; Eisert, 2013; Vervliet et al., 2014; Cichocki et al., 2015; Sidiropoulos et al., 2017; Orús, 2019; Biamonte, 2019). These properties make them well-suited for machine learning, recently leading to several promising results (Novikov et al., 2015; Stoudenmire & Schwab, 2016; Cohen et al., 2016; Glasser et al., 2019). Under some mild constraints, a TN may be interpreted as a quantum circuit (Grant et al., 2018; Huggins et al., 2019; Haghshenas et al., 2022). In some cases, these circuits may be simulated with a computational effort that scales polynomially with the number of qubits, which makes them amenable to numerical exploration on classical computers—in stark contrast with the exponential scaling expected in simulation of arbitrary quantum circuits.

One of the underlying challenges in developing quantum TN methods for machine learning is to establish a competitive performance baseline. A classical algorithm, random kitchen sinks (RKS) (Rahimi & Recht, 2007; 2008a;b), offers some inspiration since it has been shown to be competitive with multilayer neural networks (May et al., 2017). The principal difference between the two approaches is that RKS, a kernel method, replaces the costly learning of low-level features in a multilayer network with random non-local features. This randomization step may be considered a special kind of feature engineering in a machine learning pipeline.

A quantum algorithm based on RKS, known as *quantum kitchen sinks* (QKS) (Wilson et al., 2018), has been shown to produce error rates that are comparable to quantum TN methods (Huggins et al., 2019) at similar qubit utilizations and minimal classical overhead. An important question is whether coherent quantum processing can provide any improvement in classification performance. The original QKS results indicated some mild improvement with coherent processing over a small number of qubits and some degradation at a large number of qubits—and no systematic guidance for how coherent processing could be leveraged. Here we attempt to shed some light into this question by combining QKS and TN circuits into a new variational ansatz, and show this combined ansatz compares favorably to both of its ancestors.

In both the classical and quantum contexts, the contest between the multilayer networks and relatively shallow (Q/R)KS models is strongly impacted by the design of the network layers and features. It is

well-known that any function can be learned by either shallow or multilayer architecture; the essential questions are (1) how efficient is the corresponding architecture and (2) how much work is required in optimizing the architecture to the task (Bengio, 2007). We assess each of these points in turn by separately studying the two modules of the combined QKS and TN protocol.

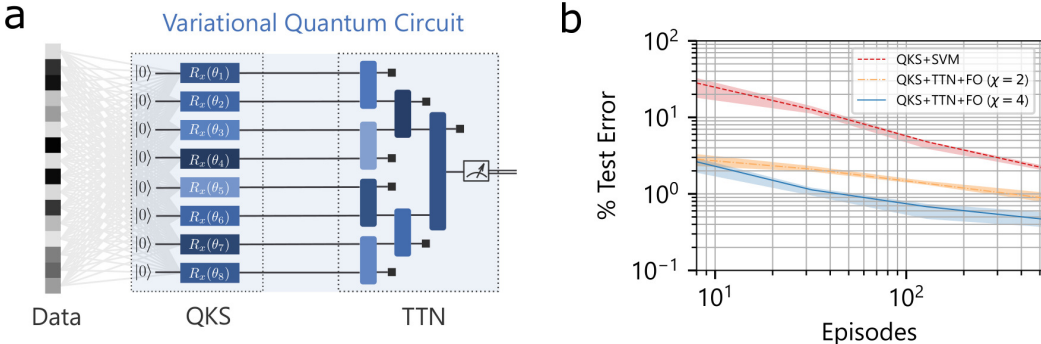


Figure 1: (a) Classical data parameterize rotations in quantum kitchen sinks (QKS). Application of quantum rotation gates such as $R_P(\theta) = \cos(\theta/2)I - i \sin(\theta/2)P$ (where P is a Pauli operator) results in quantum states whose amplitudes are mixtures of non-linear functions of the classical angles. Variational circuits where such quantum gates are followed with a tree tensor network (TTN) structure coherently process the output of QKS before a small part of the state is measured. (b) Test errors over 10 realizations (median and 68% credibility intervals) as a function of the number of episodes. Joint optimization of the rotations and TTN lead to binary classifiers with error rates below 1% (on “3” vs. “5” handwritten digits from the MNIST dataset).

A concrete example of a TN is a *tree tensor network* (TTN), depicted in (Fig. 1a). In these networks, two collections of n qubits each (where $\chi = 2^n$ is referred to as the *bond dimension*) interact unitarily, but only one of the resulting collections of n qubits continues on for additional computation. We refer to the quantum circuit that results from combining QKS with TTN as QKS+TTN, in contrast to the previous QKS approach which leveraged linear classifiers (such as support vector machines¹), which we refer to as QKS+SVM. In this work, we only considered $\chi = 2$ and $\chi = 4$ and focused on the TTN to build a large coherent computation. An attractive feature of the proposed TTN and other related ansatz (such as multi-scale entanglement renormalization ansatz (Evenbly & Vidal, 2009)) is that they avoid the problem of *barren optimization landscapes*, due to the shallow circuit depth (McClellan et al., 2018; Wang et al., 2021; Cerezo et al., 2021; Arrasmith et al., 2021), and their sparse connectivity prevents errors from accumulating pathologically (Kim & Swingle, 2017; Shehab et al., 2019; Anikeeva et al., 2021), making the implementation of these circuits on near-term hardware appealing.

Classification using QKS+TTN requires simulating the evolution of the quantum state through fixed circuits (parameterized by the classical input data), which are called kitchen sinks or *episodes* (Wilson et al., 2018), and the TTN, which concludes with the application of a measurement on the remaining n qubits at the root of the tree (Grant et al., 2018; Huggins et al., 2019). For a *single-shot* execution of the circuit, the resulting outcome corresponds to the classical output of the classifier. For a *multi-shot* execution of the circuit, we consider the application of a linear classifier to the observed outcome frequencies. For binary classification, we may consider measuring a single qubit out of the remaining n , but more general classification tasks may use all possible χ outcomes of the final measurement.

Contributions. We propose a new circuit ansatz where a TTN coherently processes the non-local feature maps of QKS. We empirically evaluate the performance of the new ansatz on image classification. In terms of classification performance, we find that simply combining QKS and TTN yields no qualitative improvements. However, training QKS+TTN with *feature optimization* (QKS+TTN+FO) significantly boosts performance and improves qubit utilization over the QKS+SVM baseline (Fig. 1b), leading to state-of-the-art quantum circuits for image classification, while requiring

¹The work of (Wilson et al., 2018) used logistic regression, while we use linear support vector machines here for convenience – we do not expect a material performance difference between the two.

only shallow quantum circuits and a small number of qubits – both well within reach of near-term quantum devices.

A crucial feature of the new ansatz is how classical data are mapped to quantum states. Previous work considered *local feature maps* (Grant et al., 2018; Huggins et al., 2019), where each classical dimension x_i was encoded in one qubit amplitude $\cos(x_i)|0\rangle + \sin(x_i)|1\rangle$ (or similar non-linear encoding), which required a number of qubits equal to the number of classical features in the problem (synthetic features may be naturally included as desired). Our work uses the *non-local feature map* provided by QKS so that the quantum state corresponding to each episode contains information about the entire data. This decouples the correlation structure in the data from the quantum circuit structure so that we may apply the same TN structure to datasets of different dimensionality and correlation structure.

2 METHODS AND BENCHMARKS

We begin by describing the datasets and the methods used for QKS+SVM, QKS+TTN, and QKS+TTN+FO. We present two approaches to reducing the complexity of the overall model: (i) translational symmetry in the TTN and (ii) sparsity in the features.

Benchmarks. The MNIST dataset (LeCun et al., 1998) is a well-known benchmark in machine learning. Each digit is a (28×28) -pixel, 8-bit grayscale image of a handwritten digit. We treat each image as a vector with dimension $28^2 = 784$ and split the dataset into 60,000 training and 10,000 testing images. While it is a benchmark for multi-class classification, we choose to focus on the binary classification two digits that are difficult to distinguish: the handwritten digits “3” and “5”, which we refer to as *(3,5)-MNIST*. All licensing information for existing assets can be found in the supplementary material.

Establishing a baseline with quantum kitchen sinks. Consider an unknown target function $f : \text{Data} \rightarrow \text{Labels}$. The general approach to “learning” the target function is to approximate it with some structured parameterized function g_θ taken from some *hypothesis space*. In deep learning, such g_θ is structured as a parameterized composition of a large number of simple non-linear functions, and the many parameters of this composition all have to be learned for optimal performance. As an alternative, Rahimi and Recht (Rahimi & Recht, 2007; 2008a;b) have shown that it was possible to produce a sufficiently rich set of hypothesis functions $g(\cdot, \theta)$ as weighted linear sums of simple non-linear functions with *random* parameters. The weights in the summary linear combination of such non-linear terms still need to be learned, but this is a linear learning step that can be done at a significantly lower cost than learning a full customary neural network. Results presented in (Rahimi & Recht, 2007; 2008a;b) suggest that there are multiple reasonable choices of non-linearities that can be used as structural blocks for RKS: cosine, sign, and indicator functions in particular.

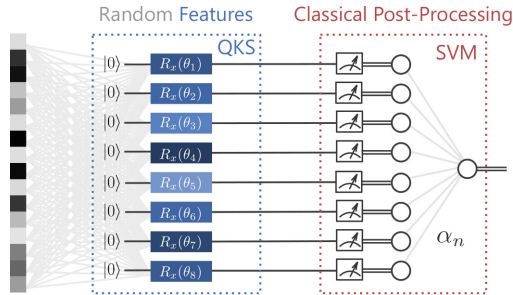


Figure 2: A quantum kitchen sink (QKS), highlighted in blue, is made up of independent episodes, each consisting of a unitary operation applied to $|0\rangle$ state. Each unitary is parameterized by a random, fixed vector (grey) and the classical input. The resulting tensor product state is then measured in a fixed basis, and the outcomes are processed further by classical post-processing (red) to generate a classical output. This procedure can be engineered to approximate any function of the classical input.

The quantum kitchen sinks (QKS) algorithm proposed by (Wilson et al., 2018) leverages the connection between trigonometric functions and qubit rotations to implement the required non-linearities.

Concretely, the original proposal used *linear* random mixtures of features as angles for single-qubit rotations, optionally followed by entangling operations. A quantum circuit composed of such operations impacts the amplitudes of a quantum state with cosines and sines of weighted mixtures of classical data.

More specifically, let $\mathbf{x} \in \mathbb{R}^p$ be a p -dimensional feature vector. Let $\mathbf{\Omega} = (\omega_1, \dots, \omega_q)^T$ be a randomized matrix, where for each k , ω_k is a p -dimensional vector with $r \leq p$ elements having random values and remaining elements being exactly zero. We can also specify a random q -dimensional bias vector β . Then we get our set of random quantum circuit parameters $\theta = \mathbf{\Omega} \mathbf{x} + \beta$.

For a sufficiently large count E interpreted as the *number of episodes* we repeat this randomized synthesis E times to form a set of *encoding parameters* $\{\mathbf{\Omega}_e, \beta_e\}_{e=1}^E$. This set of parameters is drawn only once and becomes a permanent part of a QKS solution. The $\mathbf{\Omega}_e$ and β_e elements can be drawn from various statistical distributions—e.g., normal, uniform, etc.—and parameters of these distributions become hyperparameters of the method. Choice of variances, for one, has a strong impact on the outcomes—as shown by a cross-validation grid-search in the supplementary material.

Once we have encoded the data for an episode $e \in [E]$ as vectors of circuit parameters $\theta_e(\mathbf{x}) = \mathbf{\Omega}_e \mathbf{x} + \beta_e$, $\mathbf{x} \in \text{Learningdata}$ we need to choose an ansatz for quantum circuit(s) driven by these parameters. As shown by (Wilson et al., 2018) the circuit structure can be quite simple provided that the corresponding unitary transformation depends strongly on the θ s. For example, the one-qubit ansatz presented in Fig. 2 has been proven to work well in multiple datasets. As suggested by the figure, the quantum step starts with some basic quantum state, e.g., $|0\rangle$, and the circuit is followed by a full measurement that extracts a classic bit string and collapses the state. This concludes the quantum feature pre-processing step. The bits extracted by measurement of multiple episodes form an aggregated feature vector that is then fed into a classical linear classification algorithm.

The critical difference between QKS and classical RKS is that in the case of QKS, the aggregated feature vector is by itself stochastic even though the encoding parameters $\{\mathbf{\Omega}_e, \beta_e\}_{e=1}^E$ are fixed. This is due to the stochastic nature of quantum measurement. Therefore it might be crucial to allow multiple runs (referred to as *shots*) of the same circuit within each episode and average measurement results across these runs. Such a multi-shot approach provides a more accurate representation of non-linearities induced by the quantum encoding. In either case, the totality of quantum steps generates an aggregated feature vector that is post-processed further. (Wilson et al., 2018) used Logistic Regression (LR) for the classical post-processing, while here, we explore several different approaches (both classical and quantum) to achieve better performance.

Coherently processing quantum kitchen sinks with tensor networks. Although the non-linearities provided by QKS are sufficient to approximate arbitrary functions of the classical inputs (and may therefore be used for a wide range of machine learning tasks), it is not apparent how to optimally engineer multi-qubit episodes, much less how performance may depend on the number of qubits in an episode. The results of (Wilson et al., 2018) indicated some potential improvement with the number of qubits per episode, but also considered only fixed instances of arbitrarily chosen multi-qubit circuits. In this section, we explore an approach to address this question by considering coherent processing of the output of the QKS with highly structured shallow quantum circuits: tree tensor networks (TTN). We chose to focus on TTNs due to computational convenience, but expect additional gains may be possible with other TN structures.

The original QKS proposal used tensor product measurements, followed by (largely) unstructured classical post-processing of measurement outcomes. Here, we wanted to consider measurements that have an efficiently contractible TN structure (Fig. 3). This change resulted in a variational quantum circuit with multiple layers of coherent processing. The architecture design consisted of two main modules: the non-local features and TTN.

A TTN will map $E/\lg \chi$ qubits down to n qubit in $O(\lg E)$ depth by successively discarding (or rather, ignoring) half of the qubits. Note that the TTN we employ here is a slight modification of the TTN considered in condensed matter studies (Shi et al., 2006; Tagliacozzo et al., 2009; Murg et al., 2015; Nakatani & Chan, 2013). In particular, our TTN is not made up of isometries (unitaries followed by projections on some of the outputs) because that would correspond to post-selection on measurement outcomes at each tensor and an exponentially small probability of post-selection

success for the overall circuit ². Instead, in our TTN we discard some of the output qubits for each tensor in the tree, so in some sense, our TTN is “dissipative”, with the advantage that it does not require post-selection. This changes the connectivity of the network compared to the isometric case. However, it only translates to an effective increase in the bond dimension, which is why the contraction cost has asymptotic scaling proportional to χ^7 instead of the usual χ^3 for isometric TTN (see details on computational complexity in the supplementary material).

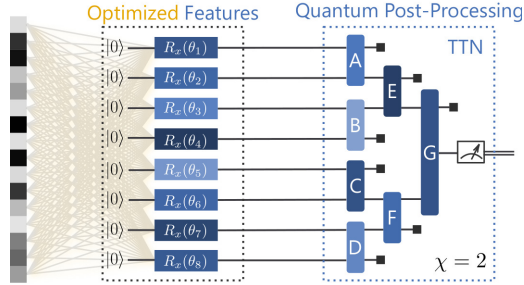


Figure 3: Coherent post-processing of the quantum features with a tree tensor network (TTN), highlighted in blue. The TTN is a variational quantum circuit consisting of unitary interactions between qubits. In the $\chi = 2$ case, a unitary—represented by a uniquely parameterized tensor—operates on two qubits. Only one of the resulting qubits continues onto the next layer, while the other is measured and discarded (as depicted by the black squares). Furthermore, in a departure from the pure randomization approach of QKS, here we also consider optimizing the features (yellow) as well as the TTN, using randomization only to initialize the QKS.

The $O(\chi^7)$ scaling can be understood by noting that the TTN can be contracted efficiently starting at the leaves. Episodes interact via $\chi^2 \times \chi^2$ unitaries (two input indices and two output indices, which are then doubled because we consider density matrices), and a partial trace must be computed to discard half of the episodes at the output of the unitary. Overall this computation requires summing over four input indices, two output indices, and a final index for the degrees of freedom that are traced out, for a total of 7 indices of dimension χ . After the leaves are contracted with the first layer of unitaries, the TTN structure is recovered, so the next layer can be contracted with the same complexity. This does require a number of contractions that is linear in E , but parallelism allows for $O(\lg E)$ time contraction.

A quantum computer may implement the contraction by decomposing the $\chi^2 \times \chi^2$ unitary into qubit interactions. This can be done in depth $O(\chi^4)$ according to the Solovay-Kitaev theorem (Nielsen et al., 2000), pointing to a potential cubic speed-up over the classical implementation. The more meaningful (and subtle) comparison to other classical classification algorithms in terms of the scaling necessary to achieve a fixed target classification accuracy is not addressed here and remains an open problem.

Architecture Design & Training. In the proposed architectural design, each input datum \mathbf{x} is translated into a product state of the form $\hat{\rho}_0(\vec{x}_n) \otimes \hat{\rho}_1(\vec{x}_n) \otimes \hat{\rho}_2(\vec{x}_n) \otimes \dots \otimes \hat{\rho}_E(\vec{x}_n)$ where E is the number of episodes. If we average over all training data in a class ℓ , we obtain the separable state $\hat{\rho}^\ell = \frac{1}{N_\ell} \sum_{n=0}^{N_\ell} \hat{\rho}_0(\vec{x}_n^\ell) \otimes \hat{\rho}_1(\vec{x}_n^\ell) \otimes \hat{\rho}_2(\vec{x}_n^\ell) \otimes \dots \otimes \hat{\rho}_E(\vec{x}_n^\ell)$ where N_ℓ is the number of training examples with class label ℓ . In the following, we assume binary classification with $\ell \in \{0, 1\}$ for simplicity. However, all the equations can be straightforwardly generalized for a multi-class setting.

Tensor network optimization. We may consider training the TTN by first fixing the QKS parameters. The unitary tensors are optimized by gradient descent methods applied to the training objective described below and in the supplementary material. Similar to (Huggins et al., 2019), the unitaries are expressed as matrix exponentials of anti-hermitian matrices, while anti-hermiticity (and thus unitarity) is preserved by an appropriate choice of parameterization.

²This may be acceptable in networks that aim to train *generative models*, which would be employed by running the networks backward and treating the measurements in the isometries as state preparation. However, such an approach cannot be directly employed with QKS.

We optimize the TTN by minimizing $\Pr(\text{error}) = 1 - \sum_{\ell \in L} \text{tr} \hat{\rho}^\ell \hat{M}_\ell / |L|$, the probability of single-shot classification error in the training set, where L is the set of labels and the \hat{M}_ℓ are the elements of a positive operator-valued measure (POVM). The \hat{M}_ℓ are implicitly defined by the TN and fixed projectors at the root of the tree TN, i.e., $\hat{M}_\ell = \mathcal{E}^\dagger(|\ell\rangle\langle\ell|)$ for some completely-positive trace-preserving (CPTP) map \mathcal{E} corresponding to the state evolution in the dissipative TN. For simplicity, we use the same objective for single-shot and multi-shot classification³, although this approach requires modification for multi-class multi-shot classification.

Due to the normalization condition on the POVM, binary classification yields $\Pr(\text{error}) = \frac{1}{2} - \frac{1}{2} \text{tr} [(\hat{\rho}_0 - \hat{\rho}_1) \mathcal{E}^\dagger(|0\rangle\langle 0|)]$, so that we may take the maximization of $f = \text{tr} [(\hat{\rho}_0 - \hat{\rho}_1) \mathcal{E}^\dagger(|0\rangle\langle 0|)]$ as our objective instead (with similar expressions for the multi-class case). It is inconvenient to manipulate $\hat{\rho}_\ell$ directly, as these matrices have size that is exponential in E , but $c_i(\bar{\ell}, \ell) = \text{tr} [\hat{\rho}_i^\ell \mathcal{E}^\dagger(|\ell\rangle\langle\ell|)]$ can be computed by contracting the TN for any given training example $\hat{\rho}_i^\ell$, which can be done on a classical computer by only manipulating tensors of a fixed dimension in $O(\chi)$ for $(\chi \geq |L|)$. On a quantum computer, $c_i(\bar{\ell}, \ell)$ may be estimated by running the state preparation and TN as circuits, and estimating the probability of obtaining outcome ℓ at the final measurement for state preparation $\hat{\rho}_i^\ell$.

Computing the objective function requires contracting the entire TN for each training example, which is excessive in some cases. This can be replaced by considering only random mini-batches of the training set using stochastic variants of gradient descent (see supplementary material). Suppose we choose to optimize each tensor sequentially. In that case, we may partially contract the TN so that at each step of the optimization, the objective reduces to the product of three small matrices—the unitary tensor, its conjugate, and a non-unitary tensor corresponding to the partial contraction of the remainder of the network which remains fixed—in a manner that is reminiscent of *quantum combs* (Chiribella et al., 2008). However, we find that globally updating all tensors leads to faster convergence and lower computational overhead.

Feature optimization. In a departure from the pure randomization approach of kitchen sinks, here we also consider optimizing the features as well as the tensor network, using randomization only to initialize the QKS. This additional optimization step can also be approached with gradient descent. For an episode $e \in [E]$, we computed the gradient with respect to the global objective f of the ansatz, $\nabla f = \left[\frac{\partial f}{\partial \omega_1}, \frac{\partial f}{\partial \omega_2}, \dots, \frac{\partial f}{\partial \omega_d}, \frac{\partial f}{\partial \beta} \right]$, with respect to parameters $\{\Omega_e, \beta_e\}$ of the driving angles $\theta_e(\mathbf{x}) = \Omega_e \mathbf{x} + \beta_e$. For a batch of examples indexed by $n \in [N_\ell]$, the gradient elements are $\frac{\partial f}{\partial \omega_i} = \frac{1}{N_\ell} \sum_n \left(\frac{\partial \rho_n}{\partial \theta_n} x_{n,i} \right) \cdot V_n$ and $\frac{\partial f}{\partial \beta} = \frac{1}{N_\ell} \sum_n \frac{\partial \rho_n}{\partial \theta_n} \cdot V_n$ where ρ_n represents the density matrix of the feature being optimized and “ $\cdot V_n$ ” represents contraction with the remainder of the tensor network. With the gradient in hand, we sequentially optimize each feature with respect to the global objective. Additional architecture and training details can be found in the supplementary material.

Reducing model complexity via translational symmetry and sparsity constraints. We explore two approaches for reducing model complexity by imposing (i) translational symmetry on the TTN (see motivation in the supplementary material) and (ii) sparsity on the features. Translational symmetry is enforced by requiring all tensors at a fixed distance from the root to be identical, while sparsity is enforced by setting a fixed (random) fraction of Ω_e to 0. These constraints reduce the size of the hypothesis space and could offer better generalization, that is, up to a point where the model begins to underfit. Parsimonious models with minimal complexity could exist right before this transition.

3 RESULTS

In this section, we present the results of several numerical experiments on QKS, QKS+TTN, and QKS+TTN+FO. We begin by establishing the baseline performance of QKS. Next, we combine QKS with a TTN—we found that simply combining QKS and TTN did not improve classification performance over the QKS baseline. However, the addition of feature optimization greatly improved the performance and yielded state-of-the-art quantum circuits for image classification. Finally, for the

³The performance in the multi-shot case is not determined directly by the objective, but rather by numerically training and testing a linear classifier on the observed outcomes frequencies for multiple experimental shots on each training/testing example (i.e., 2 dimensional real-valued vectors)

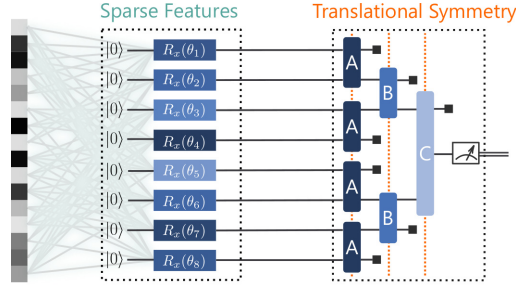


Figure 4: Translational symmetry and sparsity may be imposed on the TN and features, respectively, thereby reducing the number of parameters in the model. Note the sparser connectivity between data features and quantum rotations (teal), and the layer-wise symmetry of the tensors (orange).

most complex networks, we show that imposing translational symmetry and sparsity constraints led to dramatic reductions in the number of free parameters while having minimal impact on classification performance.

Baseline performance of quantum kitchen sinks. The original QKS design, as described by (Wilson et al., 2018), was intentionally limited to circuits where the quantum encoding of the data features was fed into a linear post-processing layer. For the (3,5)-MNIST dataset, error rates between 3.3% and 3.7% were reported in the experiments that ran on noisy quantum hardware. Results obtained on a noiseless simulator showed error rates at or slightly below 2%. While these error rates showed a lift over a linear SVMs (which was the stated intent of the original QKS research), the performance has been inferior compared to the state of the art for (3,5)-MNIST classification.

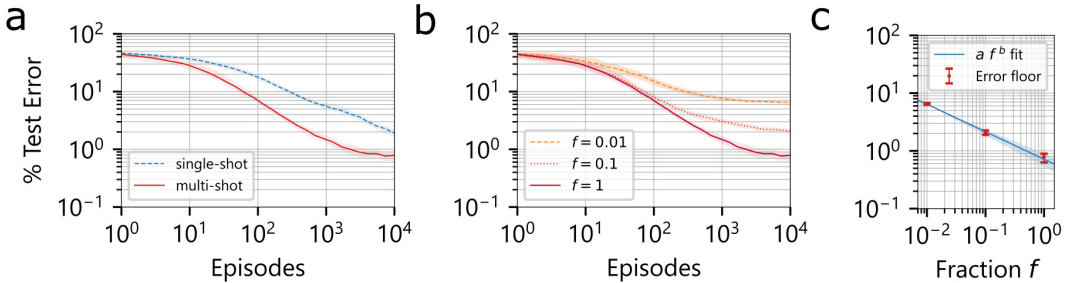


Figure 5: (a) Classification error rates for (3,5)-MNIST over 100 realizations (median and 68% credibility intervals) as a function of the number of QKS episodes. Allowing multiple shots at the one-qubit ansatz (multi-shot) led to better performance compared to single shots at the same circuits (single-shot). (b) Restricting the training dataset to a fraction of its original size ($f = 0.1$ and $f = 0.01$) limited the achievable classification performance, and (c) the observed dependency on the size of the training dataset was fit to a power law $y(f) = a f^b$ where $(a, b) = (0.72 \pm 0.18, -0.48 \pm 0.06)$, consistent with the $O(1/\sqrt{N})$ predictions from theoretical results (Rahimi & Recht, 2007; 2008a;b).

We established a more competitive baseline than the one given in the original QKS proposal by allowing multiple runs, or *shots*, of the quantum circuits. Compared to the single-shot approach, the multi-shot approach offered better error rate scaling as a function of the number of episodes (Fig. 5a). At the largest number of episodes we tested ($E = 10,000$), the single-shot approach yielded a test error of 1.87 ± 0.28 (mean \pm std.) on (3,5)-MNIST, which was consistent previous observations (Wilson et al., 2018), while the multi-shot approach offered a significant reduction to 0.78 ± 0.12 .

Further, using the multi-shot approach, we validated the scaling of the error rate as a function of the training set size (N) against theoretical guarantees for random feature algorithms. We simulated smaller datasets by taking fractions of the training dataset ($f = 0.1$ and $f = 0.01$), which shifted the noise floor higher at a large number of episodes Fig. 5b. The empirical dependence of the noise floor

on the size of the training set was fitted to a power law $y(f) = a f^b$ via a non-linear least-squares procedure yielding $(a, b) = (0.72 \pm 0.18, -0.48 \pm 0.06)$,⁴ which is consistent with the $O(1/\sqrt{N})$ predictions from theoretical results (Rahimi & Recht, 2007; 2008a;b). For the full dataset ($f = 1$) and a large number of episodes ($E \rightarrow \infty$), the power law predicts an error of $a = 0.72 \pm 0.18$, which is consistent with the error 0.78 ± 0.12 observed at $E = 10,000$. Note that the one-qubit QKS ansatz corresponds to the random Fourier features of RKS; therefore, we treat QKS+SVM as the classical model and explore improvements against this baseline via coherent post-processing.

Coherent processing of non-local features yields compact networks. From the results presented in the preceding section, it was clear that many episodes were required for QKS to perform well. Up to this point, the quantum states generated by QKS have been measured and classically post-processed with a linear SVM, and the parameters of QKS have been randomly drawn/fixed. Here we have replaced classical post-processing with quantum post-processing—QKS+SVM to QKS+TTN—and optimized the features along with the tensor network (*feature optimization*), using randomization only to initialize QKS. We found that QKS+TTN yielded no qualitative improvements, however, the addition of feature optimization (QKS+TTN+FO) greatly boosted performance.

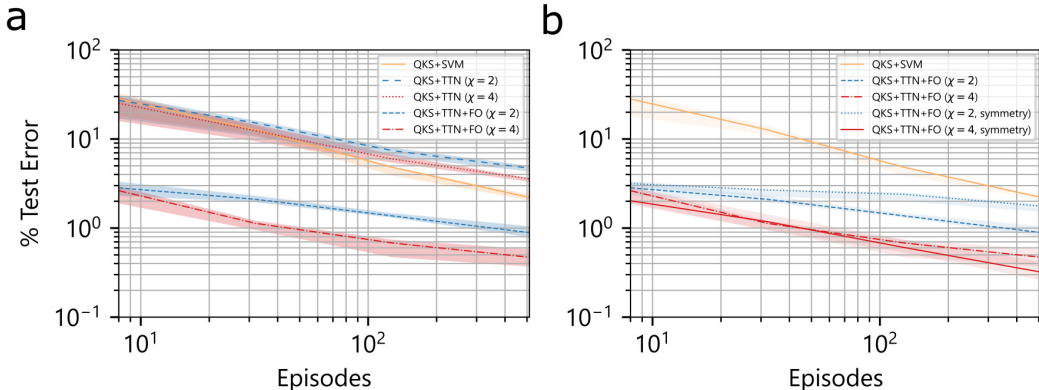


Figure 6: (a) Performance of combined protocols on (3,5)-MNIST over 10 realizations (median and 68% credibility intervals) as a function of the number of episodes or qubits. It shows that feature-optimized networks offer a significant performance lift over the QKS baseline while admitting a $20\times$ reduction in qubit utilization. (b) Imposing translational symmetry constraints on the TN degrades the performance of $\chi = 2$ networks but enhances the performance of $\chi = 4$ networks.

We compared the performance of networks *without* feature optimization (QKS+TN) against those *with* feature optimization (QKS+TTN+FO). Networks without feature optimization (QKS+TTN) closely tracked the QKS+SVM baseline while networks with feature optimization (QKS+TTN+FO) outperformed QKS+SVM, while admitting a $20\times$ reduction in the number of episodes/qubits (Fig. 6a). Therefore, it appears that *feature optimization is necessary for the formation of compact networks*.

Reducing model complexity: translational symmetry and sparsity constraints. In the previous section, we showed that QKS+TTN+FO achieved better performance than QKS+SVM at a fraction of the number of episodes. We also reduced the model complexity, as measured by the number of free (trainable) parameters by imposing translational symmetry on the TTN and sparsity on the features.

Translational symmetry constraints on the TTN had varying effects, depending on the bond dimension of the network (Fig. 6b). While performance was degraded in the $\chi = 2$ networks, the performance of $\chi = 4$ networks was equivalent or slightly better (Table 1)—while admitting an exponential reduction in the number of parameters in the TTN.

Further, in the most complex model we tested, a QKS+TTN+FO network with 512 episodes and $\chi = 4$, imposing both translational symmetry on the TTN and *sparsity* in the features revealed a regime of parsimonious models (Table 1). The design of the best performing model on (3,5)-MNIST

⁴The uncertainties are computed by linearization of the fit model in the neighborhood of the least-squares optimum, and correspond to a single standard deviation along each coordinate.

Table 1: Translational symmetry and sparsity constraints were imposed on most complex model we tested, a QKS+TTN+FO network with 512 qubits and $\chi = 4$. Imposing these constraints yielded dramatic reductions in the number of free parameters. A combination of translational symmetry and 50% sparsity constraints yielded the best test errors (mean \pm std.) on (3,5)-MNIST, in **bold**.

Model Reduction	Free Parameters (Fixed)	Test Error (%)
No symmetry or sparsity	466,945	0.44 ± 0.13
Transl. Symmetry	403,960	0.35 ± 0.07
Transl. Symmetry + Sparsity (50%)	203,256	0.35 ± 0.04
Transl. Symmetry + Sparsity (25%)	102,904	0.37 ± 0.07
Transl. Symmetry + Sparsity (13%)	52,728	0.45 ± 0.08
Transl. Symmetry + Sparsity (6%)	27,640	0.65 ± 0.18

had a combination of constraints: a TTN with translational symmetry and features with 50% sparsity. Note, however, that all parsimonious models had test errors below the QKS baseline.

State-of-the-art quantum circuits for image classification. We performed benchmarks on a broader set of classification tasks. On the hardest MNIST binary classification tasks, the non-local feature maps of QKS led to significantly lower test errors compared to local feature maps, when the pre-processing and architecture/training configurations were matched (Table 2). The best performing QKS+TTN+FO model (with translational symmetry) from the previous section led to the lowest test errors we observed. These performance gains motivated a larger investigation into the full MNIST dataset and the general multi-class problem.

Table 2: Non-local feature maps led to state-of-the-art quantum circuits for image classification, as demonstrated on the hardest MNIST binary classification problems. We ran a *control* experiment that matched the pre-processing and configuration (64 qubits, $\chi = 2$, a mini-batch size of 222, and 30 epochs) of a similar quantum tensor network approach employing local feature maps (Huggins et al., 2019). Additionally, we tested the *best* performing QKS+TTN+FO model (with translational symmetry), which involved no pre-processing and the following configuration: 512 qubits, $\chi = 4$, a mini-batch size of 32, and 40 epochs.

Quantum Encoding	3 vs. 5	4 vs. 9	7 vs. 9	3 vs. 9	2 vs. 7
Local feature maps (Huggins et al., 2019)	12.4%	12.0%	10.7%	5.9%	4.3%
Non-local feature maps (control)	5.1%	4.2%	4.5%	2.5%	2.2%
Non-local feature maps (best)	0.4%	1.1%	0.8%	0.8%	0.9%

For the multi-class classification of MNIST, we established a baseline with QKS+SVM by constructing a multi-class model composed of one-versus-one (OvO) binary classifiers: at 10,000 episodes, the multi-class test error was 2.31%. In comparison, the QKS+TTN+FO with translational symmetry, trained as OvO binary classifiers, yielded a multi-class test error of 1.16% at 512 episodes/qubits. Taken together with the complete set of classification error rates (see supplemental material), the QKS+TTN+FO offers state-of-the-art quantum circuits for image classification.

4 DISCUSSION

We have demonstrated that combining the non-local feature maps (of QKS) with coherent processing (using a TN) and feature optimization can yield significant improvements in classification error rates over the original QKS proposal, leading to state-of-the-art quantum circuits for image classification with a qubit count and circuit depth that are within reach of near-term devices.

While these gains do not translate to meaningful quantum advantage (the TTN can be efficiently contracted in a classical computer), they illustrate that variational optimization of quantum circuits can perform non-trivial tasks on relatively small devices. The tree structure of the TN considered here has also been argued by (Kim & Swingle, 2017) to have favorable noise-resilience properties, which makes them promising candidates for interesting near-term demonstrations.

REFERENCES

- Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>. Software available from tensorflow.org.
- Galit Anikeeva, Isaac H. Kim, and Patrick Hayden. Recycling qubits in near-term quantum computers. *Phys. Rev. A*, 103(4), April 2021. ISSN 2469-9926, 2469-9934. doi: 10.1103/physreva.103.042613. URL <https://doi.org/10.1103/physreva.103.042613>.
- Martin Arjovsky, Amar Shah, and Yoshua Bengio. Unitary evolution recurrent neural networks. In *International conference on machine learning*, pp. 1120–1128. PMLR, 2016.
- Andrew Arrasmith, M. Cerezo, Piotr Czarnik, Lukasz Cincio, and Patrick J. Coles. Effect of barren plateaus on gradient-free optimization. *Quantum*, 5:558, October 2021. ISSN 2521-327X. doi: 10.22331/q-2021-10-05-558. URL <https://doi.org/10.22331/q-2021-10-05-558>.
- Yoshua Bengio. On the challenge of learning complex functions. *Prog. Brain Res.*, 165:521–534, 2007.
- Yoshua Bengio, Yann LeCun, et al. Scaling learning algorithms towards AI. *Large-scale kernel machines*, 34(5):1–41, 2007.
- Debjyoti Bhattacharjee and Anupam Chattopadhyay. Depth-optimal quantum circuit placement for arbitrary topologies. *arXiv preprint arXiv:1703.08540*, 2017.
- Jacob Biamonte. Lectures on quantum tensor networks, 2019. URL <https://arxiv.org/abs/1912.10049>.
- M. Cerezo, Akira Sone, Tyler Volkoff, Lukasz Cincio, and Patrick J. Coles. Cost function dependent barren plateaus in shallow parametrized quantum circuits. *Nat. Commun.*, 12(1), March 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-21728-w. URL <https://doi.org/10.1038/s41467-021-21728-w>.
- Samuel Yen-Chi Chen, Chih-Min Huang, Chia-Wei Hsing, and Ying-Jer Kao. An end-to-end trainable hybrid classical-quantum classifier. *Machine Learning: Science and Technology*, 2(4):045021, sep 2021. doi: 10.1088/2632-2153/ac104d. URL <https://doi.org/10.1088/2632-2153/ac104d>.
- G. Chiribella, G. M. D’Ariano, and P. Perinotti. Quantum circuit architecture. *Phys. Rev. Lett.*, 101(6):060401, August 2008. ISSN 0031-9007, 1079-7114. doi: 10.1103/physrevlett.101.060401. URL <https://doi.org/10.1103/physrevlett.101.060401>.
- Andrzej Cichocki, Danilo Mandic, Lieven De Lathauwer, Guoxu Zhou, Qibin Zhao, Cesar Caiafa, and Huy Anh Phan. Tensor decompositions for signal processing applications: From two-way to multiway component analysis. *IEEE Signal Process Mag.*, 32(2):145–163, March 2015. ISSN 1053-5888. doi: 10.1109/msp.2013.2297439. URL <https://doi.org/10.1109/msp.2013.2297439>.
- N. Cohen, O. Sharir, and A. Shashua. On the expressive power of deep learning: A tensor analysis. *29th Annual Conference on Learning Theory*, (698–728), 2016. doi: <https://doi.org/10.48550/arXiv.1509.05009>. URL <https://doi.org/10.48550/arXiv.1509.05009>.
- Freeman J Dyson. The threefold way. algebraic structure of symmetry groups and ensembles in quantum mechanics. *Journal of Mathematical Physics*, 3(6):1199–1215, 1962.
- J. Eisert. Entanglement and tensor network states. 2013. doi: 10.48550/ARXIV.1308.3318. URL <https://arxiv.org/abs/1308.3318>.

- G. Evenbly and G. Vidal. Algorithms for entanglement renormalization. *Phys. Rev. B*, 79(14), April 2009. ISSN 1098-0121, 1550-235X. doi: 10.1103/physrevb.79.144108. URL <https://doi.org/10.1103/physrevb.79.144108>.
- Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors, 2018. URL <https://arxiv.org/abs/1802.06002>.
- Ivan Glasser, Ryan Sweke, Nicola Pancotti, Jens Eisert, and Ignacio Cirac. Expressive power of tensor-network factorizations for probabilistic modeling. *Adv Neural Inf Process Syst*, 32, 2019.
- Edward Grant, Marcello Benedetti, Shuxiang Cao, Andrew Hallam, Joshua Lockhart, Vid Stojevic, Andrew G. Green, and Simone Severini. Hierarchical quantum classifiers. *npj Quantum Inf.*, 4(1), December 2018. ISSN 2056-6387. doi: 10.1038/s41534-018-0116-9. URL <https://doi.org/10.1038/s41534-018-0116-9>.
- Reza Haghshenas, Johnnie Gray, Andrew C. Potter, and Garnet Kin-Lic Chan. Variational power of quantum circuit tensor networks. *Phys. Rev. X*, 12(1):011047, March 2022. ISSN 2160-3308. doi: 10.1103/physrevx.12.011047. URL <https://doi.org/10.1103/physrevx.12.011047>.
- Charles R. Harris, K. Jarrod Millman, Stéfan J. van der Walt, Ralf Gommers, Pauli Virtanen, David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J. Smith, Robert Kern, Matti Picus, Stephan Hoyer, Marten H. van Kerkwijk, Matthew Brett, Allan Haldane, Jaime Fernández del Río, Mark Wiebe, Pearu Peterson, Pierre Gérard-Marchant, Kevin Sheppard, Tyler Reddy, Warren Weckesser, Hameer Abbasi, Christoph Gohlke, and Travis E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, September 2020. doi: 10.1038/s41586-020-2649-2. URL <https://doi.org/10.1038/s41586-020-2649-2>.
- William Huggins, Piyush Patil, Bradley Mitchell, K Birgitta Whaley, and E Miles Stoudenmire. Towards quantum machine learning with tensor networks. *Quantum Sci. Technol.*, 4(2):024001, January 2019. ISSN 2058-9565. doi: 10.1088/2058-9565/aaea94. URL <https://doi.org/10.1088/2058-9565/aaea94>.
- Tak Hur, Leeseok Kim, and Daniel K Park. Quantum convolutional neural network for classical data classification. *Quantum Machine Intelligence*, 4(1):1–18, 2022.
- Li Jing, Yichen Shen, Tena Dubcek, John Peurifoy, Scott Skirlo, Yann LeCun, Max Tegmark, and Marin Soljačić. Tunable efficient unitary neural networks (eunn) and their application to rnns. In *International Conference on Machine Learning*, pp. 1733–1741. PMLR, 2017.
- Isaac H. Kim and Brian Swingle. Robust entanglement renormalization on a noisy quantum computer, 2017. URL <https://arxiv.org/abs/1711.07500>.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Tamara G. Kolda and Brett W. Bader. Tensor decompositions and applications. *SIAM Rev.*, 51(3):455–500, August 2009. ISSN 0036-1445, 1095-7200. doi: 10.1137/07070111x. URL <https://doi.org/10.1137/07070111x>.
- Y. LeCun. Generalization and network design strategies. In R. Pfeifer, Z. Schreter, F. Fogelman, and L. Steels (eds.), *Connectionism in Perspective*, Zurich, Switzerland, 1989. Elsevier. an extended version was published as a technical report of the University of Toronto.
- Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proc. IEEE*, 86(11):2278–2324, 1998. ISSN 0018-9219. doi: 10.1109/5.726791. URL <https://doi.org/10.1109/5.726791>.
- Ding Liu, Shi-Ju Ran, Peter Wittek, Cheng Peng, Raul Blázquez García, Gang Su, and Maciej Lewenstein. Machine learning by unitary tensor network of hierarchical tree structure. *New J. Phys.*, 21(7):073059, July 2019. ISSN 1367-2630. doi: 10.1088/1367-2630/ab31ef. URL <https://doi.org/10.1088/1367-2630/ab31ef>.

- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations (ICLR)*, 2019.
- Avner May, Alireza Bagheri Garakani, Zhiyun Lu, Dong Guo, Kuan Liu, Aurélien Bellet, Linxi Fan, Michael Collins, Daniel Hsu, Brian Kingsbury, Michael Picheny, and Fei Sha. Kernel approximation methods for speech recognition, 2017. URL <https://arxiv.org/abs/1701.03577>.
- Jarrod R. McClean, Sergio Boixo, Vadim N. Smelyanskiy, Ryan Babbush, and Hartmut Neven. Barren plateaus in quantum neural network training landscapes. *Nat. Commun.*, 9(1), November 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-07090-4. URL <https://doi.org/10.1038/s41467-018-07090-4>.
- V. Murg, F. Verstraete, R. Schneider, P. R. Nagy, and Ö. Legeza. Tree tensor network state with variable tensor order: An efficient multireference method for strongly correlated systems. *J. Chem. Theory Comput.*, 11(3):1027–1036, February 2015. ISSN 1549-9618, 1549-9626. doi: 10.1021/ct501187j. URL <https://doi.org/10.1021/ct501187j>.
- Naoki Nakatani and Garnet Kin-Lic Chan. Efficient tree tensor network states (TTNS) for quantum chemistry: Generalizations of the density matrix renormalization group algorithm. *J. Chem. Phys.*, 138(13):134113, April 2013. ISSN 0021-9606, 1089-7690. doi: 10.1063/1.4798639. URL <https://doi.org/10.1063/1.4798639>.
- M.A. Nielsen, I.L. Chuang, and I.L. Chuang. *Quantum Computation and Quantum Information*. Cambridge Series on Information and the Natural Sciences. Cambridge University Press, 2000. ISBN 9780521635035. URL <https://books.google.com/books?id=aai-P4V9GJ8C>.
- Alexander Novikov, Dmitrii Podoprikhin, Anton Osokin, and Dmitry P Vetrov. Tensorizing neural networks. *Adv Neural Inf Process Syst*, 28, 2015.
- Román Orús. Tensor networks for complex quantum systems. *Nature Reviews Physics*, 1(9): 538–550, August 2019. ISSN 2522-5820. doi: 10.1038/s42254-019-0086-7. URL <https://doi.org/10.1038/s42254-019-0086-7>.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Adv Neural Inf Process Syst*, 20, 2007.
- Ali Rahimi and Benjamin Recht. Uniform approximation of functions with random bases. In *2008 46th Annual Allerton Conference on Communication, Control, and Computing*, pp. 555–561. IEEE, IEEE, September 2008a. doi: 10.1109/allerton.2008.4797607. URL <https://doi.org/10.1109/allerton.2008.4797607>.
- Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. *Adv Neural Inf Process Syst*, 21, 2008b.
- Chase Roberts, Ashley Milsted, Martin Ganahl, Adam Zalcman, Bruce Fontaine, Yijian Zou, Jack Hidary, Guifre Vidal, and Stefan Leichenauer. Tensornetwork: A library for physics and machine learning, 2019.
- Matthew Rocklin. Dask: Parallel computation with blocked algorithms and task scheduling. In *Proceedings of the 14th python in science conference*, number 130-136. Citeseer, 2015.
- Omar Shehab, Isaac H. Kim, Nhung H. Nguyen, Kevin Landsman, Cinthia H. Alderete, Daiwei Zhu, C. Monroe, and Norbert M. Linke. Noise reduction using past causal cones in variational quantum algorithms, 2019. URL <https://arxiv.org/abs/1906.00476>.
- Y.-Y. Shi, L.-M. Duan, and G. Vidal. Classical simulation of quantum many-body systems with a tree tensor network. *Phys. Rev. A*, 74(2), August 2006. ISSN 1050-2947, 1094-1622. doi: 10.1103/physreva.74.022320. URL <https://doi.org/10.1103/physreva.74.022320>.

- Nicholas D. Sidiropoulos, Lieven De Lathauwer, Xiao Fu, Kejun Huang, Evangelos E. Papalexakis, and Christos Faloutsos. Tensor decomposition for signal processing and machine learning. *IEEE Trans. Signal Process.*, 65(13):3551–3582, July 2017. ISSN 1053-587X, 1941-0476. doi: 10.1109/tsp.2017.2690524. URL <https://doi.org/10.1109/tsp.2017.2690524>.
- P.Y. Simard, D. Steinkraus, and J.C. Platt. Best practices for convolutional neural networks applied to visual document analysis. In *Seventh International Conference on Document Analysis and Recognition, 2003. Proceedings.*, pp. 958–963, 2003. doi: 10.1109/ICDAR.2003.1227801.
- E Miles Stoudenmire. Learning relevant features of data with multi-scale tensor networks. *Quantum Sci. Technol.*, 3(3):034003, April 2018. ISSN 2058-9565. doi: 10.1088/2058-9565/aaba1a. URL <https://doi.org/10.1088/2058-9565/aaba1a>.
- E. Miles Stoudenmire and David J. Schwab. Supervised learning with quantum-inspired tensor networks, 2016.
- L. Tagliacozzo, G. Evenbly, and G. Vidal. Simulation of two-dimensional quantum systems using a tree tensor network that exploits the entropic area law. *Phys. Rev. B*, 80(23), December 2009. ISSN 1098-0121, 1550-235X. doi: 10.1103/physrevb.80.235127. URL <https://doi.org/10.1103/physrevb.80.235127>.
- Nico Vervliet, Otto Debals, Laurent Sorber, and Lieven De Lathauwer. Breaking the curse of dimensionality using decompositions of incomplete tensors: Tensor-based scientific computing in big data analysis. *IEEE Signal Process Mag.*, 31(5):71–79, September 2014. ISSN 1053-5888. doi: 10.1109/msp.2014.2329429. URL <https://doi.org/10.1109/msp.2014.2329429>.
- Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.
- Samson Wang, Enrico Fontana, M. Cerezo, Kunal Sharma, Akira Sone, Lukasz Cincio, and Patrick J. Coles. Noise-induced barren plateaus in variational quantum algorithms. *Nat. Commun.*, 12(1), November 2021. ISSN 2041-1723. doi: 10.1038/s41467-021-27045-6. URL <https://doi.org/10.1038/s41467-021-27045-6>.
- C. M. Wilson, J. S. Otterbach, N. Tezak, R. S. Smith, A. M. Polloreno, Peter J. Karalekas, S. Heidel, M. Sohaib Alam, G. E. Crooks, and M. P. da Silva. Quantum kitchen sinks: An algorithm for machine learning on near-term quantum computers, 2018. URL <https://arxiv.org/abs/1806.08321>.
- Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-{MNIST:} a novel image dataset for benchmarking machine learning algorithms, 2017. URL <http://arxiv.org/abs/1708.07747>.

A COMPUTATIONAL COMPLEXITY

As we have stated in the introduction, tensor network (TN) methods have been widely studied in the machine learning and physics communities due to their expressive power and computational convenience. Under some mild constraints on the components, a TN may be interpreted as quantum circuits, and in some cases, these circuits may be simulated with a computational effort that scales polynomially with the number of qubits involved—in stark contrast with the exponential scaling that may be expected for arbitrary circuits.

A concrete example of a TN that can be simulated efficiently is a *tree tensor network* (TTN). In these networks, two collections of n qubits each (where $\chi = 2^n$ is referred to as the *bond dimension*) interact unitarily⁵, but only one of the resulting collections of n qubits continues on for additional computation (the other is simply ignored). By successively discarding half of the qubits, a TTN will map $E/\lg \chi$ qubits down to n qubits in $O(\lg E)$ depth. In this work, we only considered $\chi = 2$ and $\chi = 4$, and focused on the TTN to build a large coherent computation.

Classification using QKS+TTN requires simulating the evolution of the quantum state through the fixed circuits (parameterized by the classical input data) and the TTN, which concludes with the application of a measurement on the remaining n qubits at the root of the tree. The result outcome (or its expectation value) is then used for classification.

Thanks to the tree structure of the TN, the output of QKS+TTN can be computed by a classical computer in time $O(\chi^7 E)$ if the TN is contracted serially or in time $O(\chi^7 \lg E)$ if the network is contracted in parallel. A quantum computer may implement the same transformation in time $O(\chi^3)$, with the disadvantage that the final outcome can only be sampled (but this sampling can be done in parallel by running multiple copies of the circuit). The classical data must be pre-processed for each episode, with a runtime of $O(E D)$ serially or $O(D)$ in parallel, where D is the dimensionality of the classical data.

Training the network, however, requires more than just being able to simulate the evolution of the quantum state: it requires optimizing over all the parameters in the TN. Without imposing any symmetry, the TTN has $(E-1)\chi^2(\chi^2+1)/2 \lg \chi$ real parameters. If the tensors along each layer of the TTN are constrained to be identical, the number of parameters is reduced to $\lg(E)\chi^2(\chi^2+1)/2 \lg \chi$. Overall, the cost of training will be proportional to the number of parameters in the model and to the training batch size.

The classical time overhead of classification with QKS+TTN is $O(D + \chi^7 E)$, yielding the full distributions of outcomes. The quantum time overhead with the equivalent circuits is $O(S D + S \chi^4 \lg E)$, where S is the number of runs of the quantum circuit (given we can only sample the final measurement in the quantum case, instead of obtaining the full distribution). The time overhead for the training of QKS+TTN picks up additional factors proportional to the number of parameters in the TN, but these are the same for the classical and quantum cases.

While for $\chi > 2$, the compilation of each tensor into the native gate operations may add significant classical computation overhead during training and significantly increase circuit depth, one may replace the monolithic $\chi > 2$ tensors with a network of subcomponents of bond dimension 2. Such replacement would require a number of tensors/unitaries that grows as $O(\chi^4)$ in order to achieve good accuracy in full generality (Nielsen et al., 2000), but one may consider instead the replacement with a fixed network of subcomponents, and optimizing over the parameters of that fixed network Haghshenas et al. (2022). Although we do not consider the impact of such a fixed structure here, there are indications that it would not impact classification performance—with the advantage that it would greatly reduce the number of free parameters in the network, and the TN optimization would be cheaper than computing accurate unitary decompositions at every iteration.

As pointed out in the main text, the TTNs we employ here are a slight modification of the TTNs considered in condensed matter studies. In particular, our TTN is not made up of isometries. Instead, in our TN we “ignore” or “discard” some of the output qubits for each tensor in the tree, so in some sense, our TTN is “dissipative.” This changes the connectivity of the network compared to the

⁵We also considered more general operations, such as quantum channels. Although quantum channels have an appealing structure for numerical optimization (the optimization of each local tensor is a convex problem), it did not appear to lead to any performance improvement over the strictly unitary case.

isometric case, but it only translates to an effective increase in the bond dimension, which is why the contraction cost has asymptotic scaling proportional to χ^7 instead of the usual χ^3 for isometric TTNs.

We did not take into consideration the impact of connectivity constraints of any experimental realization of a quantum computer. It should be noted, however, that connectivity constraints impose at most a linear depth overhead as a function of width (Bhattacharjee & Chattopadhyay, 2017), but more detailed (and likely architecture-dependent) analysis would be needed to pin down concrete overhead numbers.

B MATERIALS AND METHODS

B.1 ARCHITECTURE AND TRAINING DETAILS

Quantum kitchen sinks. We simulated quantum kitchen sinks (QKS) using NumPy (Harris et al., 2020) and processed the classical outputs using scikit-learn (Pedregosa et al., 2011). Hyperparameters such as the variance of QKS weights and the regularization parameter of linear SVMs were optimized using cross-validation grid search and randomized search, respectively. For each experimental condition, we collected 32 logarithmically spaced sample points and 100 realizations using Dask (Rocklin, 2015) to distribute the computation across an internal cluster with 340 nodes, where each node had 44 cores and 318 GB.

Tree tensor networks. We combined QKS with a tree tensor network (TTN) by using TensorNetwork (Roberts et al., 2019) to construct/contract the networks and SciPy (Virtanen et al., 2020) to initialize/optimize the tensors and features. The TTN was initialized by drawing N tensors from a circular unitary ensemble (CUE) (Dyson, 1962)— $U_n \sim \text{CUE}(\chi^2)$ for $n \in [N]$, where the eigenvalues of U_n have unit length and uniformly distributed in phase—then, for each unitary tensor, we computed the $\chi^2 - 1$ real numbers parameterizing the corresponding hermitian matrix H_n in $U_n = \exp(iH_n)$. To train the TTN, we chose to optimize each tensor sequentially—below, we discuss global updates to all tensors. In that case, we partially contracted the TTN so that at each step of the optimization, the objective was reduced to the product of three small matrices—the unitary tensor, its conjugate, and a non-unitary tensor corresponding to the partial contraction of the remainder of the network which remained fixed—in a manner that was reminiscent of *quantum combs* (Chiribella et al., 2008). The partial contractions were computed and summed over a batch of $B = 1024$ examples—an embarrassingly parallel computation that was distributed using Dask. This distributed operation resulted in a small matrix that could be used as a fixed input to any number of optimization methods. To optimize the corresponding unitary tensor, we used conjugate gradient (with a tolerance of 10^{-5} and a maximum number of iterations set at 100) since it led to the fastest convergence.

Feature optimization. The sequential approach to optimizing each tensor was extended to tuning each feature, which we called feature optimization (FO). The intuition behind FO was derived from classical work by Bengio and LeCun comparing kernel machines and neural networks, where the key difference was whether the low-level features (or *templates matching units*) were tunable (Bengio et al., 2007). It was argued that more compact networks could be formed by tuning the basis functions in the features to the training objective. Since the combination of QKS+TTN did not offer a performance boost over the QKS+SVM baseline, we hypothesized that compact networks could be formed by optimizing the features—i.e., QKS+TTN+FO could lead to performance gains over the QKS+SVM baseline. Note, the propagation function (or input currents) to a unit i in the the first hidden layer of a neural network $z_i = W_i x + b_i$ is equivalent to the driving angle $\theta_e(\mathbf{x}) = \Omega_e \mathbf{x} + \beta_e$ of a QKS episode. However, while the input currents are acted upon by sigmoid/rectifier activation functions in the neural network, the driving angles parameterize a qubit rotation in the case of the one-qubit QKS ansatz.

Using randomization only to initialize features, the addition of feature optimization (QKS+TTN+FO) involved several changes to the optimization of the unitary tensors. For an episode $e \in [E]$, we computed the gradient with respect to the global objective f of the ansatz,

$$\nabla f = \left[\frac{\partial f}{\partial \omega_1}, \frac{\partial f}{\partial \omega_2}, \dots, \frac{\partial f}{\partial \omega_d}, \frac{\partial f}{\partial \beta} \right],$$

with respect to parameters $\{\Omega_e, \beta_e\}$ of the driving angles $\theta_e(\mathbf{x}) = \Omega_e \mathbf{x} + \beta_e$. For a batch of examples indexed by $n \in [N_\ell]$, the gradient elements are

$$\frac{\partial f}{\partial \omega_i} = \frac{1}{N_\ell} \sum_n \left(\frac{\partial \rho_n}{\partial \theta_n} x_{n,i} \right) \cdot V_n, \quad \frac{\partial f}{\partial \beta} = \frac{1}{N_\ell} \sum_n \frac{\partial \rho_n}{\partial \theta_n} \cdot V_n,$$

where ρ_n represents the density matrix of the featuring that is being optimized and “ $\cdot V_n$ ” represents contraction with the remainder of the tensor network. Similar to the optimization of the TTN,

conjugate gradient was used for feature optimization, where the only change was that the maximum number of iterations was set at 5. For different bond dimensions ($\chi = 2$ and $\chi = 4$), we compared QKS+TTN and QKS+TTN+FO at 500 sweeps by collecting 4 linearly spaced sample points and 10 realizations using Dask on the internal cluster (described above).

Translational symmetry and sparse features. Reducing the complexity of a learning model is directly related to its generalizability and speed of convergence (LeCun, 1989). While there is a well-known 2D translational symmetry that may be exploited for image classification (Simard et al., 2003; LeCun, 1989), it was not considered in this work; rather, we focused on the permutation symmetry in the independent and identically distributed QKS episodes, with the goal of reducing the number of trainable parameters in the TTN.

Since translation symmetry is included in the permutation symmetry, we imposed layer-wise weight sharing on the TTN. Note, FO breaks the symmetry of the QKS episodes—however, reducing number of parameters by imposing translational symmetry is still well-motivated, because such a constraint is beneficial in avoiding barren plateaus (Wang et al., 2021; Cerezo et al., 2021; Arrasmith et al., 2021) and overfitting (LeCun, 1989). In short, we did not exploit the 2D translational symmetry in the images, but instead we were motivated by the permutation symmetry in the distribution of QKS episodes before FO. Imposing translational symmetry required a different training algorithm since each layer in the TTN shared the same set of parameters. We implemented a version of the architecture in TensorFlow (Abadi et al., 2015) so that we could make global updates via automatic differentiation (Algorithm 1). All parameters were updated using Adam (Kingma & Ba, 2014), similar to Grant et al. (Grant et al., 2018), a batch size of $B = 32$, and an initial learning rate of 0.001. We compared QKS+TTN+FO with and without translational symmetry in the most complex model we tested ($E = 512$ and $\chi = 4$) at 40 epochs and over 10 realizations using a single desktop with 4 cores and 32 GB.

Using the Tensorflow implementation, we also imposed sparsity in the features by setting a fixed (random) fraction of Ω_e to zero. We sought the most parsimonious QKS+TTN+FO models (for $E = 512$ and $\chi = 4$) by sampling 5 exponentially decreasing densities d at 40 epochs and over 10 realizations using the single desktop (described above).

Algorithm 1 Training algorithm for QKS+TTN+FO

```

1: given training data  $\{\mathbf{x}_m, \ell_m\}_{m=0}^M$ , batch size  $B$ , number of QKS episodes  $E$ , sparsity of features
    $d \in (0, 1]$ , bond dimension  $\chi$ , number of tensors  $N$ , and translational symmetry flag  $f_{sym}$ .
2:
3: initialize random parameters  $\Omega_e \sim \mathcal{N}(0, \sigma^2)$  and  $\beta_e \sim \mathcal{U}(0, 2\pi)$  for  $e \in [E]$ , and Haar random
   unitaries  $\{U_n\}_{n=1}^N$ , where  $U_n \sim \text{CUE}(\chi^2)$ .
4:
5:  $\Theta \leftarrow (\{\Omega_e, \beta_e\}_{e=1}^E, \{U_n\}_{n=1}^N)$ 
6:
7: function VARIATIONALCIRCUIT( $\Theta$ )
8:   for all  $e \leq E$  do
9:      $\theta_e \leftarrow \Omega_e \mathbf{x}_m + \beta_e$ 
10:     $\rho_e \leftarrow \text{PREPAREFEATURE}(\theta_e, d)$ 
11:   end for
12:    $f \leftarrow \text{CONTRACTTTN}(\{\rho_e\}_{e=0}^E, \ell_m, \chi, \{U_n\}_{n=1}^N, f_{sym})$ 
13:   return  $f$ 
14: end function
15:
16: repeat
17:    $\nabla f(\Theta) \leftarrow \text{GRADIENTONBATCH}(\text{VARIATIONALCIRCUIT}(\Theta), B)$ 
18:    $\Theta \leftarrow \text{UPDATEPARAMETERS}(\nabla f(\Theta))$ 
19: until stopping criterion is met
20: return trained classifier VARIATIONALCIRCUIT( $\Theta$ )

```

Multi-class classification. We applied the previously described architectures to a broader set of binary and multi-class classification problems by training many one-versus-one (OvO) classifiers. We tested a more general alternative that adapted the QKS+TTN+FO ($\chi = 4$) architecture for the multi-class

problem by mapping the full 16-dimensional output at the root node to 10 classes. We found that mapping 12 of the measurement outcomes to 6 classes and the remaining outcomes to 4 classes led to the fastest convergence; however, more general mappings may be considered.

All benchmarks for binary classification did not involve any pre-processing or regularization techniques. For the multi-class classification of MNIST, we added a pre-processing step that deskewed MNIST images. Additionally, to reduce overfitting and improve convergence, we replaced Adam with AdamW with warm restarts, where the learning rate and decoupled weight decay (initialized at 0.001 and 4×10^{-4}) followed cosine annealing schedules (5 restarts, $T_0 = 1$, $T_{mult} = 2$) (Loshchilov & Hutter, 2019).

B.2 LICENSING OF EXISTING ASSETS

We provide licensing information for each existing asset below:

Python Libraries

- NumPy (Harris et al., 2020) and SciPy (Virtanen et al., 2020) are BSD licensed.
- Dask (Rocklin, 2015) and scikit-learn (Pedregosa et al., 2011) are licensed under the New BSD License.
- TensorFlow (Abadi et al., 2015) and TensorNetwork (Roberts et al., 2019) are licensed under the Apache License 2.0.

Datasets

- MNIST (LeCun et al., 1998) is licensed under the Creative Commons Attribution-Share Alike 3.0 license.
- Fashion-MNIST (Xiao et al., 2017) is licensed under the MIT license.

C SUPPLEMENTARY RESULTS

C.1 HYPERPARAMETERS FOR QUANTUM KITCHEN SINKS

Several hyperparameters in quantum kitchen sinks (QKS) affected its classification performance: the variance of the normal distribution from which the weights of the random features are drawn, controlled by σ , and the number of episodes E . We performed cross-validation (5 folds) experiments using a grid search over these hyperparameters for both single-shot and multi-shot circuits (Figure C.1) and found an optimal value of $\sigma = 0.1$ led to the best validation errors at a large number of episodes ($E = 10,000$).

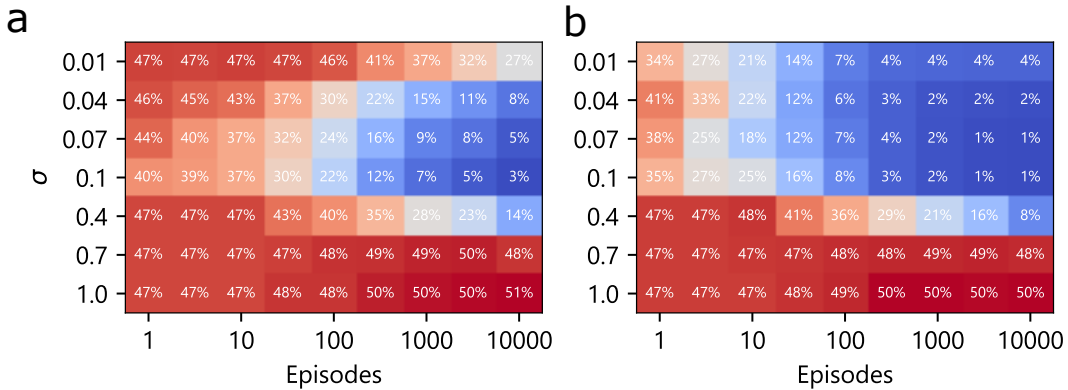


Figure C.1: (a) Single-shot validation errors (averaged over 5 folds) indicate optimal hyperparameter values: $\sigma = 0.1$ and $E = 10,000$. (b) The same hyperparameters hold for the multi-shot case.

C.2 ABLATION OF THE TENSOR NETWORK

In the main text, we showed that the addition of feature optimization (FO) to QKS+TTN was necessary for significant improvements over the QKS+SVM baseline. An important question is whether the improvements were solely due to the enrichment of the non-local features, which we refer to as *feature enrichment*. Moreover, how much of the improvement could be attributed to the *coherent post-processing* with tensor networks? We assessed the contribution of these two factors by removing the tensor network after training, and replacing the tensor network with classical post-processing (a linear SVM classifier) on the optimized non-local features (Figure C.2).

We observed different behavior for different bond dimensions. For $\chi = 2$, classical post-processing of the enriched features outperformed coherent post-processing of the enriched features, despite the features having been optimized for the coherent post-processing of the TTN. However, for $\chi = 4$, classical post-processing of the enriched features under-performed coherent post-processing of the enriched features. Interestingly, the performance of classical post-processing was essentially the same in these two cases, while the coherent processing improved by a factor of two with the larger bond dimension. This provides some indication that both feature enrichment and coherent post-processing contribute to the performance improvement of QKS+TTN+FO over QKS+TTN and QKS+SVM. In future work, this question can be more thoroughly investigated by studying feature optimization with purely classical post-processing and by exploring different tensor network features in coherent post-processing (different connectivity, higher bond dimensions, etc.).

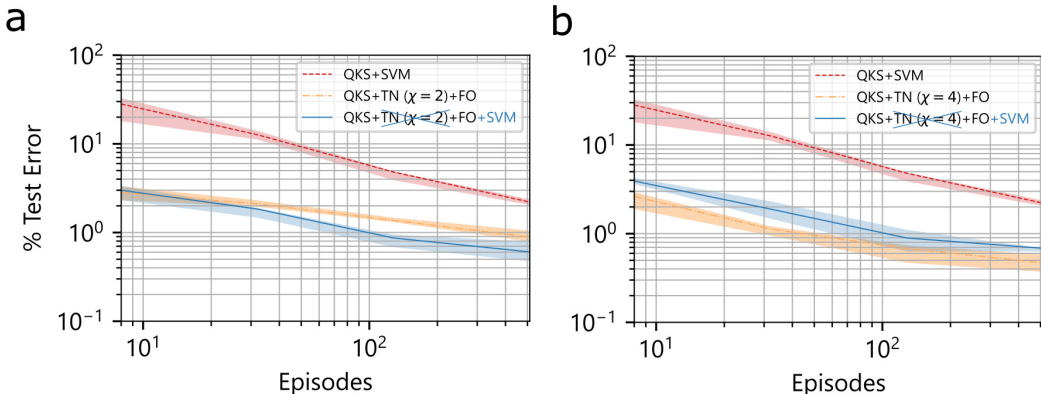


Figure C.2: (a) In the $\chi = 2$ network, the relative contribution of feature enrichment exceeded that of coherent post-processing and the reduction of the feature representation by tensor network led to a degradation in performance. (b) For the $\chi = 4$ network, there is a near-equivalent contribution of feature enrichment, however, the coherent post-processing of the tensor network offered a performance boost over classical post-processing.

C.3 (TOP, SHIRT)-FASHION-MNIST

We also considered the performance of QKS+SVM, QKS+TTN, and QKS+TTN+FO (with and without translational symmetry) on the Fashion-MNIST (Xiao et al., 2017) dataset. The Fashion-MNIST dataset (Xiao et al., 2017) is composed of 28×28 -pixel, 8-bit grayscale images, each depicting various merchandise and articles of clothing, making it a more difficult classification task. We vectorize and split the dataset similarly to that of MNIST. We tested the models on a hard binary classification problem, “Top” vs. “Shirt,” which we refer to as (*Top vs. Shirt*)-Fashion-MNIST.

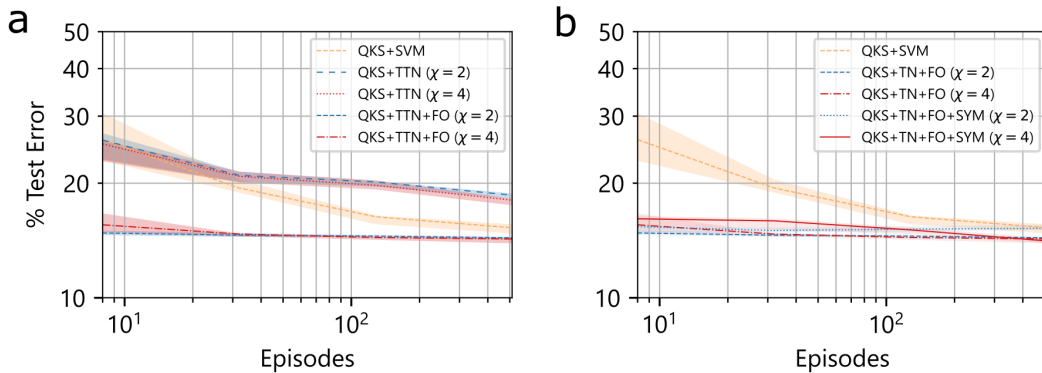


Figure C.3: (a) Classification error rates over 10 realizations (median and 68% credibility intervals) as a function of the number of episodes for (*Top vs. Shirt*)-Fashion-MNIST. Since this is a much harder classification problem, the combined QKS+TTN protocol reaches an apparent noise floor at a much smaller qubit utilization. (b) Enforcing translational symmetry constraints in the TTN degrades performance in smaller architectures but offered equivalent or better performance in the largest architecture we studied.

C.4 BENCHMARKS ON MNIST AND FASHION-MNIST

As shown in the main text, the non-local feature maps of QKS led to significantly lower test errors than local feature maps on the hardest MNIST binary classification tasks. Here we have benchmarked the best architectural configuration we tested—QKS+TTN+FO with translational symmetry, $E = 512$, and $\chi = 4$ trained for 40 epochs—on the full MNIST and Fashion-MNIST datasets (Fig. C.4). To the best of our knowledge, the model achieves state-of-the-art error rates compared to other variational quantum circuit and TN approaches (Huggins et al., 2019; Farhi & Neven, 2018; Chen et al., 2021; Hur et al., 2022).

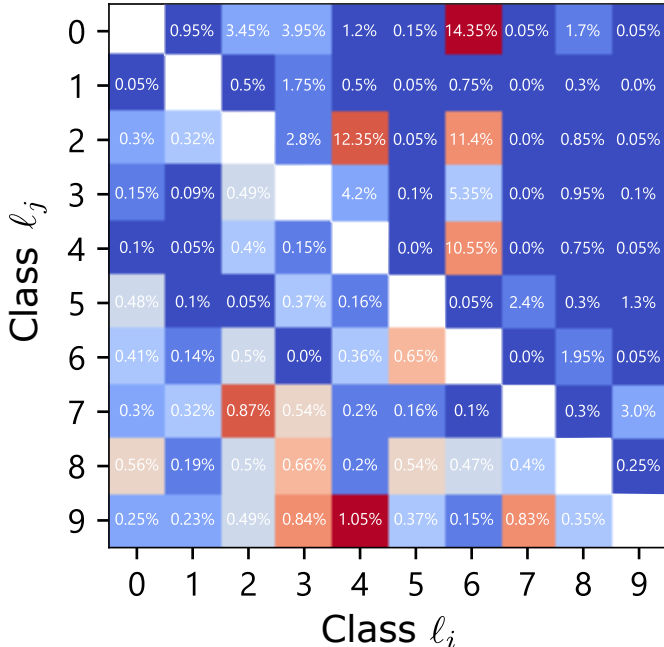


Figure C.4: Pairwise binary classification error rates for the best architectural configuration we tested—QKS+TTN+FO with translational symmetry, $E = 512$, and $\chi = 4$ trained for 40 epochs—on MNIST (lower triangle) and Fashion-MNIST (upper triangle).

To test the generalization of QKS+TTN+FO architecture, we modified it directly for the multi-class case by mapping the full 16-dimensional output at the root node to 10 classes. At 512 episodes (no symmetry), the test error reached 1.82%, while admitting a $45\times$ reduction in the number of parameters compared to the OvO construction; at 1,024 episodes, the test error reached 1.67%.

Since we described multi-class classification benchmarks for MNIST in the main text and above, here we focus on the multi-class classification of Fashion-MNIST. We established a baseline with QKS+SVM by constructing a multi-class model composed of one-versus-one (OvO) binary classifiers: at 10,000 episodes, the multi-class test error was 12.08%. In comparison, QKS+TTN+FO with translational symmetry, trained as OvO binary classifiers, yielded a multi-class test error of 11.91% at 512 episodes.

Although the classification error rates reported here are state-of-the-art for quantum circuits (Grant et al., 2018; Huggins et al., 2019; Farhi & Neven, 2018; Chen et al., 2021; Hur et al., 2022) and competitive against quantum-inspired tensor networks (Stoudenmire & Schwab, 2016; Stoudenmire, 2018; Liu et al., 2019), they are comparable to classical methods (LeCun et al., 1998). Thus, modules of the ansatz would likely benefit from further exploration of different feature maps (Simard et al., 2003) and other tensor network structures (Orús, 2019). The upshot would be a class of efficient unitary networks that could fundamentally resolve the exploding/vanishing gradient problem in deep neural networks and the detection of long-term dependencies in dynamical systems (Arjovsky et al., 2016; Jing et al., 2017).