QTML 2021

Quantum Techniques in Machine Learning

Booklet of Extended Abstracts

Wednesday November 10 2021

Quantum Advantage in Basis-Enhanced Neural Sequence Models Extended Abstract

Eric R. Anschuetz*
MIT Center for Theoretical Physics, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

Xun Gao
Department of Physics, Harvard University, 17 Oxford Street, Cambridge, MA 02138, USA

We show that under reasonable trainability assumptions, there exists an unconditional superpolynomial separation in expressivity between commonly used neural sequence models and simple quantum extensions of them. We are able to explicitly show that the sources of these separations are quantum complementarity and contextuality.

I. BACKGROUND

The use of machine learning for language translation has undergone a revolution in recent years. Historically, Bayesian networks [1] have been used for the task of language translation—most famously, *hidden Markov models* (HMMs) [2]. Though HMMs can in theory encode any sequential distribution with enough hidden units, in practice they are very inefficient. Specifically, as HMMs allow for generic transition functions, in general one can utilize a latent space only logarithmic in the size of the input for training and inference to be efficient [3].

More recently, neural sequence models such as *recurrent neural networks* (RNNs) [4–6] and *Transformers* [7] were found to be naturally much more memory efficient than HMMs by further restricting the structure of the network. Instead of allowing for generic transition functions, neural sequence models are efficiently expressible as compositions of linear transformations with a fixed nonlinearity—more generally, these models are *Lipschitz continuous*. That is, these models are differentiable almost everywhere, and almost everywhere have bounded derivative. Though these models are technically more restricted (given fixed numerical precision) when compared with HMMs, the efficiency of neural sequence models allow them to outperform HMMs on most real world sequence modeling tasks.

Simultaneously, quantum models have been studied as a possible extension of classical machine learning models, as quantum systems are believed to produce probability distributions difficult to simulate classically [8, 9]. In previous work, the authors showed an unconditional, superpolynomial advantage in the expressivity of a class of quantum HMMs over classical HMMs in translation tasks [10]. However, the restricted memory efficiency of HMMs was crucial for this separation, and thus the methods used do not extend to a superpolynomial separation on the more efficient neural network models. It was therefore unclear whether a quantum expressivity advantage held over state of the art classical methods for sequence modeling.

II. RESULTS AND IMPLICATIONS

In this work, we take advantage of the restricted structure of seq2seq models [11] (i.e. neural sequence models based on vanilla RNNs [4], LSTMs [5], etc.) and Transformers [7] to show an unconditional superpolynomial expressivity separation between them and a certain quantum extension of them. More precisely, we show that there are translation tasks that can be performed by the quantum enhanced models to zero error in perplexity, that any polynomial-equivalently sized classical seq2seq or Transformer model cannot perform to better than infinite error in perplexity. We also explicitly show that, even without guarantees on Lipschitz continuity, our previous results can be extended to *any* neural sequence-to-sequence model to give an unconditional polynomial expressivity separation.

Similar to our previous results on HMMs [10], the simplicity of the basis-enhanced models allows us to directly track the origins of the quantum advantage. In the case of seq2seq models, we show that the origin is quantum contextuality; for Transformers, quantum complementarity. The quantum advantage therefore in some sense stems from using quantum resources to combat the inherent weaknesses of the underlying classical models. Classical seq2seq models struggle with representing the context of words in long sequences [5], and Transformer models (even with a positional encoding) are constrained to be approximately permutation invariant in order to efficiently capture long-range correlations [7].

* eans@mit.edu

First, we define the quantum extension we use to show this separation. We extend the notion of basis-enhanced Bayesian networks [10] to neural networks:

Definition 1 (Basis-enhanced neural networks)**.** A quantum circuit is a basis-enhanced neural network if it extends an implementation of a neural function with measurement allowed in any local basis.

Note that this is a very simple (and nonuniversal) quantum generalization—given an implementation of a classical neural network on a quantum computer, our model is only extended by allowing for a final layer of single qubit rotations. We then consider the sequence modeling task:

Definition 2 (Stabilizer sequence translation, informal)**.** Let $\boldsymbol{s}, \boldsymbol{s'}$ each be binary representations of a sequence of k Pauli matrices and ± 1-valued measurement outcomes on n qubits. The *stabilizer sequence translation task* is the modeling of the conditional distribution $p\left(\boldsymbol{s'} \mid \boldsymbol{s}\right)$, which is uniform over all *correct translations* $\boldsymbol{s'}$ of $\boldsymbol{s}$. A translation $\boldsymbol{s'}$ of $\boldsymbol{s}$ is considered correct if the sequential measurement scenario $(\boldsymbol{s}, \boldsymbol{s'})$ when beginning in the all zero state is consistent with quantum mechanics.

This sequence modeling task is just a classical description of a stabilizer measurement scenario. Due to the relative simplicity in which quantum computers can sample from this distribution—simply simulate the measurement scenario and then sample—it is straightforward to show that the task given in Definition 2 can be modeled completely by basis-enhanced neural networks. That is:

Theorem 1 (Basis-enhanced neural networks can stabilizer sequence translate, informal)**.** *Basis-enhanced seq2seq models with* $\mathrm{O}\left(n\right)$*-dimensional latent spaces can perform the translation task of Definition 2 to zero error.*

Using techniques from [10], it is also fairly straightforward to show that *no* (autoregressive) classical neural sequence model with a $\mathrm{o}\left(n^2\right)$-dimensional latent space can perform this task to a finite perplexity; intuitively, this stems from a memory lower bound on the classical simulation of stabilizer measurements on n qubits [12]. However, it is less straightforward to show—and what we spend the majority of our work showing—that this can be extended to a superpolynomial advantage for Lipschitz continuous (i.e. efficiently trainable) neural sequence models. Indeed, we show this separation for both seq2seq and Transformer models, with the only constraints being Lipschitz continuity of the models, the models having fixed numerical precision, and a simple trainability assumption on the seq2seq model.

Theorem 2 (Neural networks cannot stabilizer sequence translate, informal)**.** *Neither seq2seq nor Transformer models (that are Lipschitz continuous and at fixed precision) with* $O\left(\mathrm{poly}\left(n\right)\right)$*-dimensional latent spaces can perform the translation task of Definition 2 to finite error in forward or backward perplexity.*

Our proof techniques vary for the two classes of models. For seq2seq models, we show that Lipschitz continuity greatly constrains the amount of information flow through the model. This constrained information flow makes it provably difficult for seq2seq models to learn the measurement context of an observable in the translation task we consider, leading to incorrect measurement outcomes and thus an inability to perform the translation task of Definition 2 to finite error in perplexity. For Transformers, we show that the approximate permutation symmetry of certain layers of the model—an approximate symmetry that is necessary for the model to efficiently capture long-range correlations [13]—makes certain sequences difficult for the model to distinguish. In particular, we show that there must exist two measurement sequences describing orthogonal states that the model struggles to differentiate, leading to incorrect measurement outcomes as in the seq2seq case.

Our results are the first to show an unconditional superpolynomial expressivity advantage between state of the art neural networks and a quantum extension of them. Furthermore, we show that this separation arises from very basic quantum phenomena—contextuality and complementarity—and that only simple quantum models are necessary for the separation. Importantly, this provides an avenue for demonstrating these advantages experimentally; as generic quantum models are expected to suffer from difficulties in training due to features such as barren plateaus [14] in the training landscape, simpler and more physically motivated quantum models are more promising to implement. Finally, certain restricted classes of basis-enhanced models are classically simulable, giving physical motivation from quantum contextuality and complementarity that may help in designing improved classical neural networks in the future.

---

[1] J. Pearl, in *Proceedings of the 7th Conference of the Cognitive Science Society, University of California, Irvine, CA, USA* (1985) pp. 329–334.

[2] L. E. Baum and T. Petrie, The Annals of Mathematical Statistics **37**, 1554 (1966).

[3] L. E. Baum, T. Petrie, G. Soules, and N. Weiss, Ann. Math. Statist. **41**, 164 (1970).
[4] J. J. Hopfield, Proceedings of the National Academy of Sciences **79**, 2554 (1982).
[5] S. Hochreiter and J. Schmidhuber, Neural Comput. **9**, 1735 (1997).
[6] K. Cho, B. van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, and Y. Bengio, in *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)* (Association for Computational Linguistics, Doha, Qatar, 2014) pp. 1724–1734.
[7] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, in *Proceedings of the 31st International Conference on Neural Information Processing Systems*, NIPS'17 (Curran Associates Inc., Red Hook, NY, USA, 2017) pp. 6000–6010.
[8] R. P. Feynman, Int. J. Theor. Phys. **21**, 467 (1982).
[9] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, *et al.*, Nature **574**, 505 (2019).
[10] X. Gao, E. R. Anschuetz, S.-T. Wang, J. I. Cirac, and M. D. Lukin, Enhancing generative models via quantum correlations (2021), arXiv:2101.08354 [quant-ph].
[11] I. Sutskever, O. Vinyals, and Q. V. Le, in *Advances in Neural Information Processing Systems*, Vol. 27, edited by Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014).
[12] A. Karanjai, J. J. Wallman, and S. D. Bartlett, Contextuality bounds the efficiency of classical simulation of quantum processes (2018), arXiv:1802.07744 [quant-ph].
[13] T. Luong, H. Pham, and C. D. Manning, in *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing* (Association for Computational Linguistics, Lisbon, Portugal, 2015) pp. 1412–1421.
[14] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Nat. Commun. **9**, 4812 (2018).

Kernel Matrix Completion for Offline Quantum-Enhanced Machine Learning

Annie Naveh[1], Anna Phan[2,3], Imogen Fitzgerald[1], Andrew Lockwood[1], and Travis L. Scholten[4]

[1] *Woodside Energy Ltd., Perth, WA, Australia*
[2] *IBM Quantum, Melbourne, Victoria, Australia*
[3] *School of Physics, University of Melbourne, Parkville, Victoria, Australia*
[4] *IBM Quantum, IBM T.J. Watson Research Center, Yorktown Heights, New York, USA*

Abstract

Hybrid quantum-classical machine learning workflows involving quantum kernel matrices incur data transfer costs scaling quadratically with the data set size. This work shows matrix completion algorithms mitigate these costs while being robust to shot noise. The relationship between properties of quantum circuits and completability of quantum kernel matrices is studied.

Enhancing classical machine learning algorithms through quantum kernels (similarity measures) is a rapidly-growing research area in quantum machine learning [7, 9, 10, 8]. Given a data set $\{\mathbf{d}_j\}_{j=1}^N$, with $\mathbf{d}_j \in \mathbb{R}^M$, the corresponding quantum kernel matrix K has elements $K_{lm} = \mathrm{Tr}(\rho(\mathbf{d}_l)\rho(\mathbf{d}_m))$, where $\rho(\mathbf{d}_j) = U^\dagger(\mathbf{d}_j)\rho_0 U(\mathbf{d}_j)$ for a parameterized quantum circuit (PQC) $U(\mathbf{z})$ and fiducial state ρ_0. K is positive semi-definite (PSD) (i.e., $K \geq 0$), and can be used in various kernel-based classical machine learning algorithms such as Kernel Ridge Regression, Support Vector Regression, and Gaussian Process Regression.

One highly practical (yet under-appreciated) aspect of utilizing quantum kernels in such algorithms is that most workflows involve acquiring *new* data points, necessitating an extension of the quantum kernel matrix. When k new data points are acquired, $\mathcal{O}(k(N+k))$ new kernel values must be calculated. The time required to extend K in this way may exceed timescales relevant for using the machine learning algorithm.

Suppose that instead of calculating *all* the new elements of K, only *some* are. Then, classical *matrix completion* [2] can be used to fill in the remainder, thereby reducing the number of new elements that need to be computed every time a new batch of data is acquired. To formalize this, suppose K is an $N \times N$ quantum kernel matrix whose *known* elements are indexed by a set S: the tuple $(l, m) \in S$ if, and only if, K_{lm} is known. Completing K means estimating $K_{l'm'}\ \forall\ (l', m') \notin S$, yielding an estimate $\hat{K}$ of K.

The completion of K with respect to S is non-unique, and therefore is cast as an optimization problem. We choose a completion based on maximizing $\log(\det(\hat{K}))$, as it makes the fewest assumptions about K itself. The choice of the completion algorithm is guided by practical considerations of real-world workflows. First, the algorithm must run "offline" with respect to the quantum computer, meaning the elements $K_{lm}\ \forall\ (l, m) \in S$ must be calculated upfront. This mitigates increased costs inherent in adaptive, "online" completions. Second, the algorithm should *deterministically* select S to minimize the amount of data transfer, without compromising completion accuracy.

A suite of PSD matrix completion techniques based on graph theory satisfy these desiderata [3]. One of them, based on inducing a *block-diagonal sparsity pattern* [5] on the known elements, emerges as the best candidate. In this sparsity pattern, the known elements of K form a block-diagonal pattern, with overlap $u \geq 0$ between the blocks. If $u \geq r$, where $r = \mathrm{rank}(K)$, then K can be completed with zero error [4].

We numerically study completing quantum kernel matrices using the above algorithm and sparsity pattern. To do so, we generate realizations of quantum kernel matrices (using the PQC denoted "Circuit 3" in [6]) using a uniformly distributed data set to demonstrate the extension of a 450×450 matrix to include 50 new items. The 500×500 incomplete matrix has a block-diagonal sparsity pattern consisting of two blocks. The first is 450×450 (the original matrix), and the second has a size which is varied from 50×50 to 500×500. The overlap u between them is thus in the range $0 \leq u \leq 450$. For each realization, we generate an estimate $\hat{K}$, and quantify the estimation error as $\mathrm{Error} = (\|K_{\bar{S}} - \hat{K}_{\bar{S}}\|_F)/\|K_{\bar{S}}\|_F$, where $\|.\|_F$ is the Frobenius norm, and $\bar{S}$ is the complement of S with respect to the indices of K.

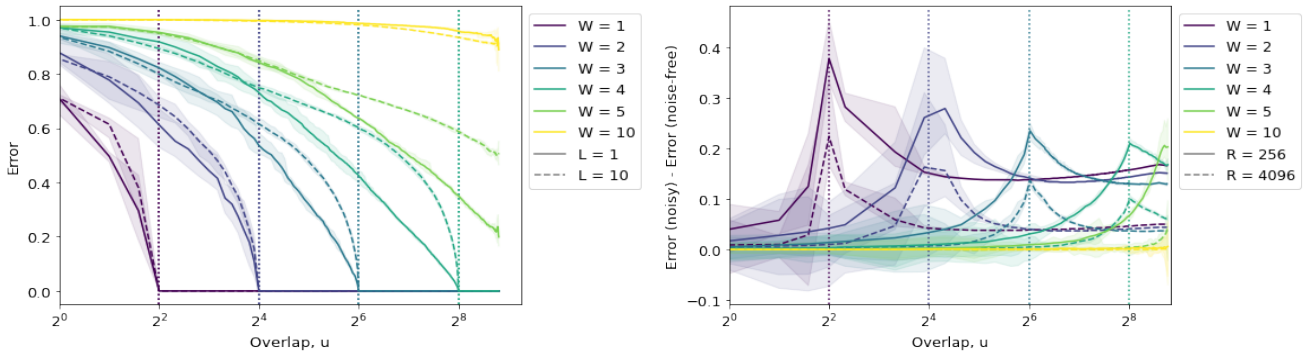

(a) **Noise-free case.** (b) **Noisy (shot noise) case.**

Figure 1: **Completing quantum kernel matrices using classical matrix completion techniques.** W denotes the width of the PQC, and L the number of repetitions of its base template. Dotted lines denote the rank r of the noise-free matrix. **Figure 1a** shows completion error goes to 0 once the overlap between the blocks u exceeds r, as predicted by [4]. **Figure 1b** shows the effect of finite-sampling ("shot") noise (R) on completion error, here L=1. In both figures, the error shading represents the minimum and maximum error from five realizations.

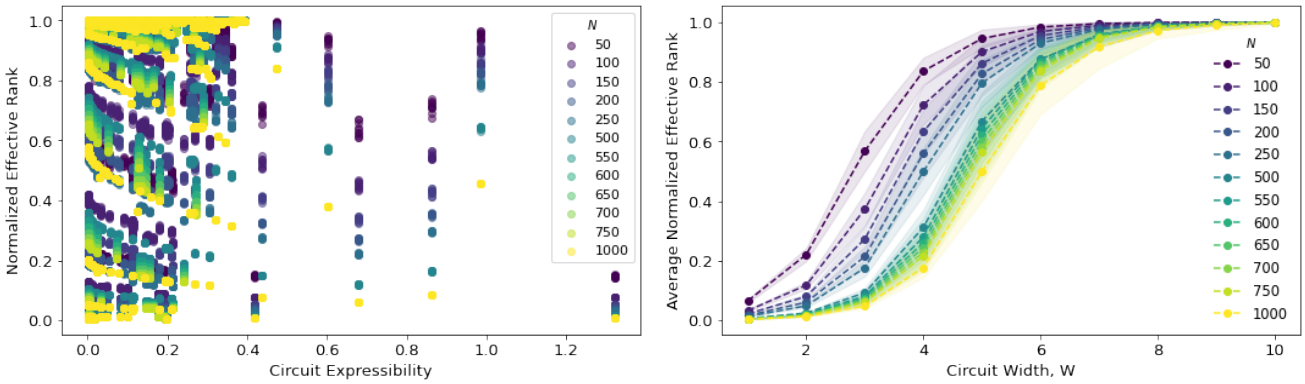

(a) **Circuit expressibility does not predict effective rank.** (b) **Circuit width predicts average effective rank.**

Figure 2: **Relationship between properties of PQCs and normalized effective rank.** For each of the 19 circuits in [6], we generate quantum kernel matrices of varying sizes (N) and compute the effective rank, normalized to N. **Figure 2a** shows circuit expressibility [6] does not seem to be predictive of normalized effective rank, whereas **Figure 2b** shows a relationship exists between normalized effective rank and circuit width.

Our first set of findings, shown in Figure 1, are (a) noise-free quantum kernel matrices can successfully be completed when $u \geq r$ (Figure 1a), and (b) completion error degrades gracefully in the presence of finite-sampling ("shot") noise (Figure 1b).

Our second set of findings, shown in Figure 2, is that different properties of PQCs may be useful for predicting a particular property of quantum kernel matrices; namely, their *effective* rank. The effective rank of a matrix [1] is a continuous measure of matrix rank, and also bounds the matrix rank from below. Predicting this number *a priori* would be useful for estimating completion error given a PQC. Figure 2a shows that circuit *expressibility* [6] does not seem to be predictive of the effective rank (normalized to matrix size) of quantum kernel matrices. However, Figure 2b shows circuit width does provide some predictive capability.

In sum, this work shows a graph-theory-based matrix completion algorithm can successfully complete quantum kernel matrices, and confirms completion error behaves as predicted. However, the quantum kernel matrices used here were generated from a random, unstructured data set, which is not wholly reflective of the structured data expected in a real-world data set. This raises an important question - how does the *structure* of a data set, through the action of a PQC, affect the (effective) rank of the quantum kernel matrix it generates? We will answer this question by applying the techniques presented here to such a data set. Further, we plan to train a kernel-based regression model with completed quantum kernel matrices to test the stability of the model under matrix completion.

References

[1] Olivier Roy and Martin Vetterli. "The effective rank: A measure of effective dimensionality". In: *2007 15th European Signal Processing Conference*. 2007, pp. 606–610.

[2] Emmanuel J Candès and Benjamin Recht. "Exact matrix completion via convex optimization". In: *Foundations of Computational mathematics* 9.6 (2009), pp. 717–772.

[3] Martin S Andersen, Joachim Dahl, and Lieven Vandenberghe. "Logarithmic barriers for sparse matrix cones". In: *Optimization Methods and Software* 28.3 (2013), pp. 396–423.

[4] William E Bishop and Byron M Yu. "Deterministic symmetric positive semidefinite matrix completion". In: *Advances in Neural Information Processing Systems* 27 (2014), pp. 2762–2770.

[5] Lieven Vandenberghe and Martin S Andersen. "Chordal graphs and semidefinite optimization". In: *Foundations and Trends in Optimization* 1.4 (2015), pp. 241–433.

[6] Sukin Sim, Peter D Johnson, and Alán Aspuru-Guzik. "Expressibility and entangling capability of parameterized quantum circuits for hybrid quantum-classical algorithms". In: *Advanced Quantum Technologies* 2.12 (2019), p. 1900070.

[7] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. "A rigorous and robust quantum speed-up in supervised machine learning". In: *arXiv preprint arXiv:2010.02174* (2020).

[8] Jennifer R Glick et al. "Covariant quantum kernels for data with group structure". In: *arXiv preprint arXiv:2105.03406* (2021).

[9] Hsin-Yuan Huang et al. "Power of data in quantum machine learning". In: *Nature communications* 12.1 (2021), pp. 1–9.

[10] Maria Schuld. "Supervised quantum machine learning models are kernel methods". In: *arXiv preprint arXiv:2101.11020* (2021).

Quantum-inspired manifold learning

Akshat Kumar[1], Mohan Sarovar[2]

[1] *Department of Mathematics, Clarkson University, Potsdam, NY 13699 USA*
[2] *Extreme-Scale Data Science and Analytics, Sandia National Laboratories, Livermore, California 94550 USA*

We introduce a (classical) algorithm for extracting geodesic distances on sampled manifolds that relies on simulation of quantum dynamics on a graph embedding of the sampled data. Our approach exploits classic results in the quantum-classical correspondence, and reveals interesting connections between the discretization provided by sampling and quantization.

Physical processes have inspired information processing algorithms throughout the history of computing. Perhaps the most prominent examples are sampling algorithms inspired by statistical mechanics such as the Metropolis-Hastings algorithm. More recently, random walks and diffusion have provided inspiration for techniques that learn the geometry of data sets and perform nonlinear dimensionality reduction to reduce the complexity of high-dimensional datasets. This task is referred to as *manifold learning*, and has become a cornerstone task in the unsupervised learning of datasets. Underlying manifold learning is the *manifold hypothesis*, which states that most real-world high-dimensional datasets, especially those originating from physical systems constrained by physical laws, are actually constrained to lower dimensional manifolds.

We introduce an algorithm for extracting geodesic distances on sampled manifolds that relies on properties of quantum dynamics and the quantum-classical correspondence. It forms a basis for techniques to learn the manifold from which a dataset is sampled, and subsequently for nonlinear dimensional reduction of high-dimensional datasets. The ingredients to our approach are (i) the approximation of a unitary time evolution operator (propagator) for a quantization of Hamiltonian dynamics on the data manifold, (ii) propagation of localized coherent states that approximate classical point particle trajectories, (iii) a rescaling of the diffusion approximation parameter to support the quantized dynamics. Quantization of the dynamics linearizes the classically non-linear problem of solving for geodesics on a manifold, and moreover, our construction allows us to use results in semiclassical analysis to prove strong asymptotic connections between the dynamics induced by unitary propagator and geodesic distances on the underlying manifold.

The key steps of quantum-inspired manifold learning are shown in Fig. 1(a). The input to the procedure is a point cloud dataset: $V = \{v_1, v_2, ..., v_N\}$, which are N samples from an underlying compact, smooth, and boundaryless Riemannian manifold $\mathcal{M}$ with dimension ν and metric tensor g_{ij}. The manifold is assumed to be isometrically embedded in $\mathbb{R}^n$ with $n > \nu$ and each data point $v_\ell \in \mathbb{R}^n$ specifies a point in these *extrinsic coordinates*. Then step 1 in the approach is to construct a data-driven approximation of a unitary time propagator for a quantum system whose configuration space is the data manifold. This is achieved by exploiting established results on convergence of normalized graph Laplacians based on sampled data to the Laplace-Beltrami operator on a data manifold. Then in step 2, we propagate data-driven approximations of minimum uncertainty coherent states on the manifold using the propagator constructed in step 1. We prove that as these states are propagated for short times, they remain localized along the geodesic flow on the manifold using a classic result from quantum-classical correspondence, Egorov's theorem, which is graphically summarized in Fig. 1(b). This enables extraction of geodesic distances from the propagated states in step 3, which is used as input for various applications (ovals on the right).

Our constructions and approach to manifold learning reveal an interesting connection between the discretization provided by finite sampling of a space and the concept of quantization. We connect the phase space uncertainty factor, h, which is a constant in physical theories, to the resolution of the manifold provided by the finite number of samples N and a scale parameter

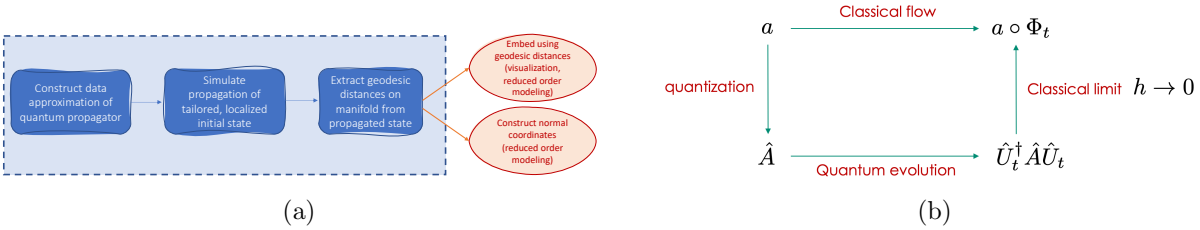

FIG. 1: **(a)** The key steps of quantum-inspired manifold learning. **(b)** Graphical summary of Egorov's theorem, which states that the classical flow of an observable a can be recovered in the $h \to 0$ limit of quantum evolution of a corresponding quantum observable, $\hat{A}$, under certain conditions.

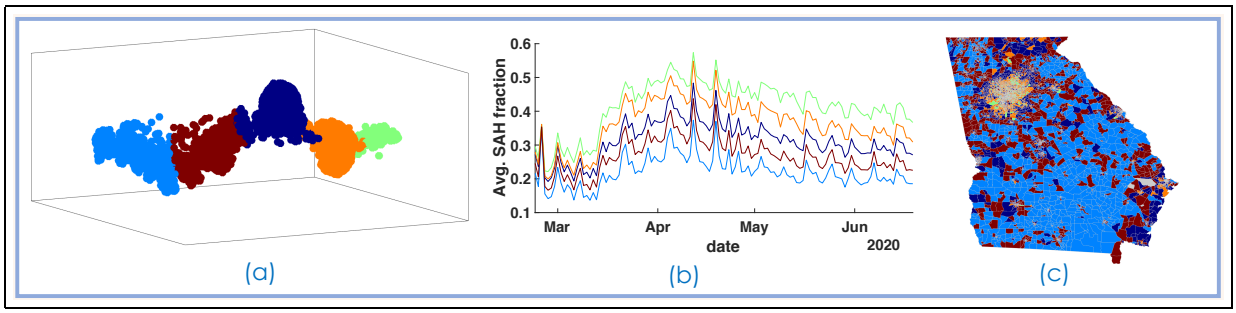

FIG. 2: An application of geodesic extraction through quantum dynamics on mobility data.

$\epsilon(N)$, which governs the quality of the approximation of free propagation of coherent states on the manifold. This $h(\epsilon(N))$ relation is intended to set the quantization scale of the quantum system to the resolution of the manifold provided by the sampled data V. With our choice of coherent states as initial states, we redistribute this resolution equally between the position and momentum degrees of freedom. From this perspective, the $h \to 0$ limit is both the classical limit of the quantum dynamics and also the limit where the sampled data fully covers the manifold.

Fig. 2 presents an application of our method to a real-world dataset, and demonstrates its utility for visualization and clustering of high-dimensional data. We analyze aggregated mobility data in the Social Distancing Metric dataset from SafeGraph Inc. This is a fine-grained dataset that collects geolocation information from mobile devices, aggregates it at the census block level (CBG) level. We extract a simple metric to gauge the mobility patterns of citizens: a daily stay-at-home (SAH) fraction for a CBG that provides a measure of how curtailed mobility was within the CBG. We performed quantum-inspired manifold learning on mobility data for the state of Georgia (GA), and after extracting geodesic distances we embed the data in three dimensions using a force-directed layout. The data is then clustered into 5 clusters using k-means. Fig. 2(a) shows the embedding in 3D with the clusters indicated by colors. In addition, the average SAH fraction time series for each cluster in Fig. 2(b) shows that the clustering is meaningful; *i.e.,* although almost all CBGs exhibit similar mobility patterns, and the clusters reflect different amounts of overall SAH fraction. Finally, Fig. 2(c) shows a map of the CBGs, which clearly shows the rural-urban divide in degree of mobility change during the pandemic; the higher SAH fraction clusters are associated with the urban centers in GA, while the majority of CBGs in the state exhibited behavior consistent with lower SAH fraction curves (light blue and brown clusters).

This work was supported by the Laboratory Directed Research and Development program at Sandia National Laboratories, a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. SAND2021-10773 A.

Multi-armed quantum bandits: Exploration versus exploitation when learning properties of quantum states

Josep Lumbreras[1], Erkka Haapasalo[1], Marco Tomamichel[1,2]

[1] *Centre for Quantum Technologies, National University of Singapore, Singapore and*

[2] *Department of Electrical and Computer Engineering,*

Faculty of Engineering, National University of Singapore, Singapore

We initiate the study of exploration/exploitation trade-offs in online learning of properties of quantum states giving an extension of the multi-armed stochastic bandit problem to an online quantum learning scenario. We quantify the difficulty of the learning problem with the minimax regret and provide various lower and upper bounds.

Note: A technical version of this work can be found at arXiv:2108.13050.

Quantum algorithms for the classical multi-armed stochastic bandit problem have been proposed recently [1, 2]. A quantum version of the Hedging algorithm, which is related to the adversarial bandit model, has also been studied [3]. These algorithms investigate potential improvements on the respective classical bandit algorithms when a quantum learner is given superposition access to the oracle, i.e., it can probe rewards for several arms in superposition. Our work generalizes multi-armed bandits from a different perspective, focusing instead on the learning theory for quantum states. In our model, which we call the multi-armed quantum bandit model, the arms correspond to different observables or measurements, the environment is an unknown quantum state and the rewards are given by the measurements on the unknown quantum state and distributed according to Born's rule. We are interested in the optimal tradeoff between exploration (i.e. learning more about the unknown quantum state) and exploitation (i.e. using acquired information about the unknown state to choose the most rewarding but not necessarily most informative measurement).

The model. A d-dimensional multi-armed quantum bandit is given by a set of observables $\mathcal{A}$ that we call actions. The bandit lives in an environment, a quantum state ρ, that is unknown but taken from a subset Γ of potential environments. Given access to n copies of the unknown quantum state, a learner at each round $t \in \{1, ..., n\}$, chooses an observable $O_t \in \mathcal{A}$, performs a measurement on ρ and receives a reward X_t, the outcome of the measurement. The learner is characterized by a conditional probability distribution over the set of actions $\mathcal{A}$ at each time step t that we denote $\pi_t(O_t|O_1, X_1, ..., O_{t-1}, X_{t-1})$ and is called *policy* or *algorithm*. The policy uses the previous rewards and observables measured in order to decide the next observable to pick from the set of actions $\mathcal{A}$. The goal of the learner will be to maximize his reward at each round.

The regret. In order to quantify the performance of the policy, we propose as a figure of merit, the cumulative expected regret, defined as

$$R_n(\mathcal{A}, \rho, \pi) := \sum_{t=1}^{n} \max_{O_i \in \mathcal{A}} \mathrm{Tr}(\rho O_i) - \mathbb{E}_{\rho,\pi}[\mathrm{Tr}(\rho O_t)]$$

where ρ is the unknown quantum state, the sum is over all rounds the learner interacts with the environment, O_t is the choice of observable by the learner at round t, and the expectation value is taken over the probability measure induced by the policy in ρ. In order to determine how well a policy performs for a given set of density matrices Γ, we define the *worst-case regret* as $R_n(\mathcal{A}, \Gamma, \pi) = \sup_{\rho\in\Gamma} R_n(\mathcal{A}, \rho, \pi)$. Moreover, in order to determine how difficult a learning problem is for a given set Γ of environments in a worst-case scenario, the *minimax regret* is defined as $R_n(\mathcal{A}, \Gamma) = \inf_{\pi\in\mathcal{P}} R_n(\mathcal{A}, \Gamma, \pi)$, where $\mathcal{P}$ is the set of all policies. A small value of $R_n(\mathcal{A}, \Gamma)$ means that the learning problem is less difficult.

The main goal of our work is to give bounds on the minimax regret in terms of the number of rounds n and the dimension of the Hilbert space d for different sets of environments and actions. Specifically, we will focus on mixed states ($\Gamma = \mathcal{S}_d$) and pure states ($\Gamma = \mathcal{S}_d^*$). For the lower bounds we use information theoretical techniques, focusing on finding environments that perform "bad" for all policies. For the upper bounds we use an algorithmic approach, i.e, try to find an algorithm that minimizes the regret and perform the regret analysis.

Results. We summarize our main results in Table I. For the case of mixed states environments and discrete action sets our lower and upper bounds match. For the case of continuous action sets containing all rank-1 projectors our upper and lower bound match in terms of the number of rounds n but there is a gap in the dimension of the system d^2. For the case of pure states environments and action set $\mathcal{A}$ containing all rank-1 projectors the regret can be expressed as $R_n = \sum_{t=1}^{n} \left(\frac{1}{2}\|\rho - \Pi_t\|_1\right)^2$ where $\rho = |\psi\rangle\langle\psi|$ is the unknown pure state and $\Pi_t \in \mathcal{A}$ the projector at time step t. This regret now has a resemblance with the regret in online learning of quantum states since it is a natural loss function for our estimate Π_t of the state ρ. We are not able to provide a non-trivial lower bound for this setting and pose this as an open question. However for this case we consider an alternative version of the regret called trace distance regret and defined as $\widetilde{R}_n(\mathcal{A}, \rho, \pi) := \frac{1}{2}\sum_{t=1}^{n} \mathbb{E}_{\rho,\pi} \|\rho - \Pi_t\|$. The lower and upper bounds for this setting with trace distance regret does not match but suggest that we can find a better algorithm such that the trace distance regret matches the lower bound. Also we provide support to the conjecture that the standard regret for this case should grow slower than $\mathcal{O}(\sqrt{n})$ if we find a suitable algorithm.

<table>
<tr><th rowspan="2"></th><th colspan="3">Discrete</th><th rowspan="2">Continous
(all rank-1 projectors)</th></tr>
<tr><th>Generic</th><th>Arm-limited</th><th>Dimension-limited</th></tr>
<tr><td rowspan="2">$\Gamma = S_d$</td><td>$R_n(\mathcal{A}, \Gamma) = \Omega(\sqrt{n})$</td><td>$R_n(\mathcal{A}, \Gamma) = \Omega(\sqrt{kn})$ $(k < d^2)$</td><td>$R_n(\mathcal{A}, \Gamma) = \Omega\left(d\sqrt{n}\right)$</td><td>$R_n(\mathcal{A}, \Gamma) = \Omega\left(\sqrt{n}\right)$</td></tr>
<tr><td rowspan="2">$R_n(\mathcal{A}, \Gamma) = \widetilde{O}(\sqrt{n})$</td><td rowspan="2">$R_n(\mathcal{A}, \Gamma) = \widetilde{O}(\sqrt{kn})$</td><td rowspan="2">$R_n(\mathcal{A}, \Gamma) = \widetilde{O}\left(d\sqrt{n}\right)$</td><td>$R_n(\mathcal{A}, \Gamma) = \widetilde{O}(d^2\sqrt{n})$</td></tr>
<tr><td rowspan="2">$\Gamma = S_d^*$</td><td>$\widetilde{R}_n(\mathcal{A}, \Gamma) = \widetilde{O}(n^{\frac{2}{3}}), (d = 2)$</td></tr>
<tr><td colspan="3">$R_n(\mathcal{A}, \Gamma) = \Omega\left(\sqrt{n}\right)$ $(d = 2, k = 3)$</td><td>$\widetilde{R}_n(\mathcal{A}, \Gamma) = \Omega\left(\sqrt{n}\right)$</td></tr>
</table>

TABLE I. Scaling of the minimax regret in terms of the number of rounds, n, dimension of the Hilbert space, d, and number of actions, k (for discrete action sets). We differentiate between arm-limited action sets (the number of arms is smaller than the degrees of freedom in the quantum state space) and dimension-limited (the number of arms can be arbitrary large).

[1] B. Casalé, G. Di Molfetta, H. Kadri, et al. Quantum bandits. *Quantum Mach. Intell.*, 2, 11, 2020.

[2] D. Wang, X. You, T. Li and A. Childs. Quantum exploration algorithms for multi-armed bandits. quant-ph/07049, 2020.

[3] P. Rebentrost *et al.* Quantum algorithms for hedging and the learning of Ising models. *Phys. Rev. A*, 103, 012418, 2021.

Provable superior accuracy in machine learned quantum models

Chengran Yang[1,2,6*], Andrew J. P. Garner[1,2,3], Feiyang Liu,[4],
Nora Tischler[5], Jayne Thompson[6], Man-Hong Yung[4],
Mile Gu[1,2,6*], Oscar Dahlsten[4,7*]

Abstract – In building models from big data, there is a constant trade-off between model precision and complexity. Here we show that quantum models enable a more favourable trade-off through machine learning techniques. We develop an algorithm to infer dimensionally constrained quantum models whose accuracy we prove exceeds any classical counterpart. We refer to arxiv: 2105.14434 for more details.

Introduction – In data analytics, the curse of dimensionality is a well-acquainted adversary [1]. As we seek to make predictions from time-series data drawn from processes of ever-growing complexity, modelling the possible future effects from all possible past observations becomes quickly intractable. Thus, when it comes to modelling highly non-Markovian processes, significant distortions become unavoidable. Here we ask: *Can dimensionally constrained quantum models – quantum machines that record relevant past information using qubits – reduce such distortions?* To answer this, we develop a quantum model discovery algorithm that infers such dimensionally constrained models directly from raw data.

Classical models and distortion – Consider a physical system observed at discrete times t with corresponding outcomes $x_t \in \mathcal{X}$ over some finite alphabet $\mathcal{X}$. We assume that the system is stationary and stochastic, such that each x_t is drawn from random variable X_t, all of which are governed by a time-translation invariant distribution $P(\overleftarrow{X}, \overrightarrow{X})$. Here, $\overleftarrow{X} := \cdots X_{-2}X_{-1}$ and $\overrightarrow{X} := X_0X_1X_2\cdots$ denote random variables governing stochastic outcomes in the past and future.

An (exact) *predictive model* is a systematic algorithm that takes each possible past $\overleftarrow{x}$, encodes it via the deterministic map $\mathcal{E}$ to some suitable state $\mathcal{E}(\overleftarrow{x})$ within a physical memory M, such that the same systematic action $\mathcal{P}$ on M at each subsequent time-step sequentially outputs $x_0, x_1, x_2\cdots$ with probability $P(\overrightarrow{X} = \overrightarrow{x}|\overleftarrow{x})$. Knowing the state of the memory M is thus as useful as $\overleftarrow{x}$ for purposes of inferring future statistics. The number of causal states, d_c, is known as a process's *topological complexity* – capturing the minimal number of configurations a predictive model requires, whilst exhibiting correct conditional future behaviour. In such scenarios of diverging d_c, finite memory means a necessary *distortion* of predictive accuracy. Our goal is then to find a model such that the deviation of $\hat{P}$ from P– the distortion–is minimal.

Quantum Models – A d-dimensional quantum model is one where each past $\overleftarrow{x}$ is encoded into a suitable quantum state $|\sigma_{\overleftarrow{x}}\rangle$ within a d- dimensional quantum memory M. Meanwhile, there must exist some quantum process whose repeated application allows sampling of the appropriate conditional futures. At each time-step, M is coupled via a time-step independent unitary U to an output register initially in $|0\rangle$. The output register is then measured, resulting in emitted outcomes $x_1, x_2, \ldots$ governed by $\hat{P}(\overrightarrow{X}|\sigma_{\overleftarrow{x}})$. There are processes where quantum models can achieve zero distortion (i.e. $\hat{P}(\overrightarrow{X}|\sigma_{\overleftarrow{x}}) = P(\overrightarrow{X}|\overleftarrow{x})$) with a lower memory dimension than classically possible [2].

Quantum model discovery algorithm – (see Box 1). The algorithm takes two inputs: (i) a data sequence $x_{0:L}$, assumed to be drawn from some stationary stochastic process P and (ii) our desired memory dimension $\hat{d}_q$. From this data, the algorithm learns a quantum model of the process, including (i) an encod-

[1]Nanyang Quantum Hub, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore 637371. [2]Complexity Institute, Nanyang Technological University, Singapore 637335. [3]Institute for Quantum Optics and Quantum Information,Austrian Academy of Sciences, Boltzmanngasse 3, Vienna 1090, Austria. [4]Shenzhen Institute for Quantum Science and Engineering and Department of Physics, Southern University of Science and Technology, Shenzhen 518055, China. [5]Centre for Quantum Computation and Communication Technology, Australian Research Council, Centre for Quantum Dynamics, Griffith University, Brisbane, Queensland 4111, Australia. [6] Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, Singapore 117543. [7]Wolfson College, University of Oxford, Oxford, OX2 6UD, United Kingdom.

Box 1: Quantum model discovery algorithm.

Algorithm:

1. To learn the unitary operator U:
   (a) Randomly initialize parameter set A corresponding up to completeness to the set of Kraus operator matrices $\{A_x := \langle x|U|0\rangle\}$.
   (b) Evaluate the cost function $C := -\log P(x_{0:L}; A)$ and its gradient ∇C.
   (c) Update the Kraus operators $\{A_x\}$ using gradient-descent based method (such as Adam).
   (d) Repeat (b)-(c) until the cost function decrease is sufficiently small.
   (e) Save final A as candidate parameters.
   (f) Repeat (a)-(e) as desired, to minimize the impact of initial choice of A. Choose the final A with the lowest C.
   (g) For that A, recover the completeness relation for $\{A_x\}$.
   (h) Construct a unitary operator U from these $\{A_x\}$.
2. To compute the encoding $\mathcal{E}$:
   (a) Compute the leading eigenvector $|\sigma_0\rangle$ of one of the Kraus operators (A_0).
   (b) $\mathcal{E}$ is then defined as $\mathcal{E}(\overleftarrow{x}) := A_{\overleftarrow{x}}|\sigma_0\rangle / \|A_{\overleftarrow{x}}|\sigma_0\rangle\|_2$.

ing map $\mathcal{E}$ which specifies how to encode each past into a quantum memory of dimension $\hat{d}_q$, and (ii) the relevant physical process, whose repeated application produces entangled output states. Measurements on such output states emit outputs that approximate the process's conditional future behaviour.

Provable Quantum Advantage – While sophisticated numerical methods for inferring classical models exist [3, 4], they do not tell us whether the resulting models are provably optimal in minimizing KL-divergence. We establish provable upper bounds on the quality of classical methods of fixed dimension. We will show that our inferred quantum models can exceed these optimal bounds.

We now apply our inference algorithm to an iconic non-Markovian process - the family of discrete-time renewal processes. Renewal processes represent an interpolation between Poissonian processes and ticking clocks, characterising many natural processes including models of lifetimes [5] and neural spike trains [6]. At each time-step, the process may emit one of two outputs: 0, representing no tick, and 1 representing a tick, and their probability of ticking at each time-step depends only on the number of time-steps since their previous tick [7–9].

We applied our algorithm to infer memory-constrained models for a variety of memory limits. We then compare the error with a lower bound of optimal classical models with similar memory constraints. Our results show that quantum models outperform classical counterparts away from the zero-error regime. To examine the resilience of the quantum advantage on present-day noisy devices, we implement the quantum model on one of IBM's quantum systems ("ibmq athens") for the case using a single qubit. Such quantum models still exhibit statistically significant accuracy advantage over the classical bound.

Conclusion – Our work substantially expands upon prior work in memory-enhanced quantum models, which focus on the memory advantages in the exact modelling of a given stochastic process [2, 9, 10]. The results here illustrate that when the memory dimension is fixed, quantum models can exhibit a significant accuracy advantage. This establishes the advantage of quantum models in real-world settings, where there is a trade-off between model complexity and accuracy. In this context, our results provide a quantum counterpart to classical tools for structural inference [4, 11, 12] and dimensional reduction in stochastic modelling [13].

[1] Verleysen, M. & François, D. The curse of dimensionality in data mining and time series prediction. In *International work-conference on artificial neural networks*, 758–770 (Springer, 2005).

[2] Thompson, J. *et al.* Causal Asymmetry in a Quantum World. *Physical Review X* **8**, 031013 (2018). URL http://arxiv.org/abs/1712.02368. 1712.02368.

[3] Kelly, D., Dillingham, M., Hudson, A. & Wiesner, K. A new method for inferring hidden markov models from noisy time sequences. *PloS one* **7**, e29703 (2012).

[4] Shalizi, C. & Klinkner, K. L. Blind construction of optimal nonlinear recursive predictors for discrete sequences. *arXiv preprint arXiv:1408.2025* (2014).

[5] Doob, J. L. Renewal theory from the point of view of the theory of probability. *Transactions of the American Mathematical Society* **63**, 422–438 (1948).

[6] Crutchfield, J. P., DeWeese, M. R. & Marzen, S. E. Time resolution dependence of information measures for spiking neurons: Scaling and universality. *Frontiers in computational neuroscience* **9**, 105 (2015).

[7] Elliott, T. J. & Gu, M. Superior memory efficiency of quantum devices for the simulation of continuous-time stochastic processes. *npj Quantum Information* **4**, 18 (2018). URL http://www.nature.com/articles/s41534-018-0064-4.

[8] Marzen, S. & Crutchfield, J. P. Informational and Causal Architecture of Continuous-time Renewal Processes. *Journal of Statistical Physics* **168**, 109–127 (2017). URL https://doi.org/10.1007/s10955-017-1793-z.

[9] Elliott, T. J. *et al.* Extreme dimensionality reduction with quantum modeling. *Phys. Rev. Lett.* **125**, 260501 (2020). URL https://link.aps.org/doi/10.1103/PhysRevLett.125.260501.

[10] Ghafari, F. *et al.* Dimensional quantum memory advantage in the simulation of stochastic processes. *Phys. Rev. X* **9**, 041013 (2019). URL https://link.aps.org/doi/10.1103/PhysRevX.9.041013.

[11] Shalizi, C. R., Shalizi, K. L. & Crutchfield, J. P. An algorithm for pattern discovery in time series. *arXiv preprint cs/0210025* (2002).

[12] Hinton, G. E., Sejnowski, T. J., Poggio, T. A. *et al. Unsupervised learning: foundations of neural computation* (MIT press, 1999).

[13] van der Maaten, L., Postma, E. & van den Herik, J. Dimensionality Reduction: A Comparative Review. *Technical Report, Tilburg University* (2007). URL https://lvdmaaten.github.io/publications/papers/TR{_}Dimensionality{_}Reduction{_}Review{_}2009.pdf.

Parametrized quantum circuits for reinforcement learning

Sofiene Jerbi, Andrea Skolik, Casper Gyurik, Simon C. Marshall, Hans J. Briegel, Vedran Dunjko

Abstract

We propose two hybrid quantum-classical reinforcement learning models, which we show can be effectively trained to solve several standard benchmarking environments. Moreover, we demonstrate and formally prove the ability of parametrized quantum circuits to solve certain learning tasks that are intractable to classical models, under widely-believed complexity theoretic assumptions.

Deep neural networks have had a profound impact on the field of reinforcement learning by recently achieving unprecedented performance in challenging decision-making tasks [1–4]. Almost in parallel, in the world of near term quantum computers, the idea that limited quantum computations, called parametrized quantum circuits [5–7] or quantum neural networks [8], could be used as building blocks of hybrid quantum-classical machine learning systems started gaining increasing traction. Such hybrid systems have already shown the potential to tackle real-world tasks in supervised [9, 10] and generative learning [11, 12], and recent works have established their provable advantages in special artificial tasks [6, 13–15]. Yet, in the case of reinforcement learning, which is arguably most challenging and where learning boosts would be extremely valuable, no proposal has been successful in solving even standard benchmarking tasks, nor in showing a theoretical learning advantage over classical algorithms. In this work, we achieve both. We find numerically that shallow quantum circuits acting on very few qubits are competitive with deep neural networks on well-established benchmarking environments. Moreover, we demonstrate, and formally prove, the ability of parametrized quantum circuits to solve certain learning problems that classical models, including deep neural networks, cannot, under the widely-believed classical hardness of the discrete logarithm problem. This constitutes clear evidence of the power of quantum learning agents and suggests that important reinforcement learning applications such as robotics [16], biology [17], or healthcare [18], can be meaningfully impacted by quantum machine learning.

This abstract summarizes two of our works, both available on arXiv [19, 20].

References

[1] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[2] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fidjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015.

[3] Piotr Mirowski, Matt Grimes, Mateusz Malinowski, Karl Moritz Hermann, Keith Anderson, Denis Teplyashin, Karen Simonyan, Andrew Zisserman, Raia Hadsell, et al. Learning to navigate in cities without a map. *Advances in Neural Information Processing Systems*, 31:2419–2430, 2018.

[4] Max Jaderberg, Wojciech M Czarnecki, Iain Dunning, Luke Marris, Guy Lever, Antonio Garcia Castaneda, Charles Beattie, Neil C Rabinowitz, Ari S Morcos, Avraham Ruderman, et al. Human-level performance in 3d multiplayer games with population-based reinforcement learning. *Science*, 364(6443):859–865, 2019.

[5] Marcello Benedetti, Erika Lloyd, Stefan Sack, and Mattia Fiorentini. Parameterized quantum circuits as machine learning models. *Quantum Science and Technology*, 4(4):043001, 2019.

[6] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567(7747):209–212, 2019.

[7] Maria Schuld and Nathan Killoran. Quantum machine learning in feature hilbert spaces. *Physical review letters*, 122(4):040504, 2019.

[8] Edward Farhi and Hartmut Neven. Classification with quantum neural networks on near term processors. *arXiv preprint arXiv:1802.06002*, 2018.

[9] Maria Schuld, Alex Bocharov, Krysta M Svore, and Nathan Wiebe. Circuit-centric quantum classifiers. *Physical Review A*, 101(3):032308, 2020.

[10] Evan Peters, Joao Caldeira, Alan Ho, Stefan Leichenauer, Masoud Mohseni, Hartmut Neven, Panagiotis Spentzouris, Doug Strain, and Gabriel N Perdue. Machine learning of high dimensional data on a noisy quantum processor. *arXiv preprint arXiv:2101.09581*, 2021.

[11] Daiwei Zhu, Norbert M Linke, Marcello Benedetti, Kevin A Landsman, Nhung H Nguyen, C Huerta Alderete, Alejandro Perdomo-Ortiz, Nathan Korda, A Garfoot, Charles Brecque, et al. Training of quantum circuits on a hybrid quantum computer. *Science advances*, 5(10):eaaw9918, 2019.

[12] Jin-Guo Liu and Lei Wang. Differentiable learning of quantum circuit born machines. *Physical Review A*, 98(6):062324, 2018.

[13] Yunchao Liu, Srinivasan Arunachalam, and Kristan Temme. A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, pages 1–5, 2021.

[14] Hsin-Yuan Huang, Michael Broughton, Masoud Mohseni, Ryan Babbush, Sergio Boixo, Hartmut Neven, and Jarrod R McClean. Power of data in quantum machine learning. *Nature communications*, 12(1):1–9, 2021.

[15] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert. On the quantum versus classical learnability of discrete distributions. *Quantum*, 5:417, 2021.

[16] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

[17] Mufti Mahmud, Mohammed Shamim Kaiser, Amir Hussain, and Stefano Vassanelli. Applications of deep learning and reinforcement learning to biological data. *IEEE transactions on neural networks and learning systems*, 29(6):2063–2079, 2018.

[18] Chao Yu, Jiming Liu, and Shamim Nemati. Reinforcement learning in healthcare: A survey. *arXiv preprint arXiv:1908.08796*, 2019.

[19] Sofiene Jerbi, Casper Gyurik, Simon Marshall, Hans J Briegel, and Vedran Dunjko. Variational quantum policies for reinforcement learning. *arXiv preprint arXiv:2103.05577*, 2021.

[20] Andrea Skolik, Sofiene Jerbi, and Vedran Dunjko. Quantum agents in the gym: a variational quantum algorithm for deep q-learning. *arXiv preprint arXiv:2103.15084*, 2021.

Structural risk minimization for quantum linear classifiers

Casper Gyurik, Dyon van Vreumingen, and Vedran Dunjko

Abstract

We provide new insights into optimally tuning parameterized quantum circuit based machine learning models by balancing the model's complexity against its success at fitting trainingdata (called structural risk minimization). Specifically, we theoretically quantify how its expressivity and empirical performance depend on the rank and Frobenius norm of the observables measured.

Link to full paper: https://arxiv.org/abs/2105.05566

Quantum machine learning (QML) models based on parameterized quantum circuits are often highlighted as candidates for quantum computing's near-term "killer application". However, the understanding of the models' ability to capture correlations in training data, and its ability to generalize to unseen data is still in its infancy. In our work we investigate these topics for a widely studied family of QML models based on parameterized quantum circuits, which includes the two prominent models introduced by Havlíček et al. [1], and Schuld and Killoran [2]. We use tools from statistical learning theory to better understand the empirical and generalization performance of these models. In particular, our objective is to study how to tune certain parameters of the QML model to balance between its training accuracy and generalization performance – a principle which is often referred to as *structural risk minimization* – in order to achieve the optimal performance in practice. To do so, we investigate two widely utilized *complexity measures* – i.e., the VC dimension and fat-shattering dimension – of the QML models, which closely captures the generalization performance of these models. In particular, we establish upper bounds on these complexity measures that explicitly depend on certain parameters of the QML model, by exploiting their relationship to linear classifiers (a well understood family of classical models). Finally, by utilizing this explicit dependence of the upper bounds on the model parameters, we devise methods that balance the empirical and generalization performance of these models, enabling the models to perform better in practice.

Let us more precisely introduce the QML model that we study. First, datapoints x are encoded into quantum states $\rho(x)$. This can for instance be achieved using a parameterized quantum circuit U via the mapping $x \mapsto \rho(x) = |\Phi(x)\rangle\langle\Phi(x)|$, where $|\Phi(x)\rangle = U(x)|0\rangle$. In our work we will not be concerned about the details of this encoding – which is extensively studied elsewhere – but we draw our attention towards optimally tuning the subsequent step in the model. After having encoded the data into a quantum state, an observable $\mathcal{O}$ chosen from a predefined family of observables $\mathbb{O}$ is measured, and depending on whether the expectation value lies above of below a threshold d a label $+1$ or -1 is assigned. In short, the family of quantum classifiers is given by

$$\mathcal{C}(\mathbb{O}) = \Big\{ c_{\mathcal{O},d}(x) = \mathrm{sign}\big(\mathrm{Tr}\big[\mathcal{O}\rho(x)\big] - d\big) \;\Big|\; \mathcal{O} \in \mathbb{O},\ d \in \mathbb{R} \Big\}.$$

This includes the two models introduced by Havlíček et al. [1], and Schuld and Killoran [2], as the family $\mathbb{O}$ can correspond to the observables implementable using a parameterized quantum circuit followed by measurement in the computational basis and postprocessing of the outcome*, or the observables that lie in the span of $\rho(x)$ for training examples x†.

*Called the quantum variational classifier [1], or the explicit approach [2].

†Called the quantum kernel estimator [1], or the implicit approach [2].

For our first result, we show that the VC dimension of the family of quantum classifiers depends on the dimension of the sum of the images of the observables. Specifically, we prove the following proposition.

Proposition 1. $\mathrm{VC}\big(\mathcal{C}(\mathbb{O})\big) \leq r^2 + 1$, *where* $r = \dim\big(\sum_{\mathcal{O}\in\mathbb{O}} \mathrm{Im}\,\mathcal{O}\big)$[§].

While this bound it not useful for arbitrary ansatze used in the explicit approach (also called the quantum variational classifier), we can design ansatzes that allow us to directly control the upper bound on the VC dimension by varying the rank of the measurement after the parameterized quantum circuit. In particular, we come up with ansatze for which we can control the upper bound on the VC dimension – and thus the generalization performance of the model – by varying the number of computational basis states upon which the final measurement projects.

For our second result, we show that the fat-shattering dimension of the family of quantum classifiers depends on the Frobenius norm of the observables. Specifically, we prove the following proposition.

Proposition 2. $\mathrm{fat}_{\mathcal{C}(\mathbb{O})}(\gamma) \leq O\left(\frac{\eta^2}{\gamma^2}\right)$, *where* $\eta = \max_{\mathcal{O}\in\mathbb{O}} \|\mathcal{O}\|_F$.

In particular, the above proposition establishes that the fat-shattering dimension – and thus the generalization performance of the model – can be controlled by varying the Frobenius norm of the observables.

In our remaining results, we study the influence that the above model parameters (i.e., the quantity r in Proposition 1 and η in Proposition 2) have on the empirical performance of the model. First, we show that quantum models that use high-rank observable can achieve strictly smaller training errors than quantum models that use low-rank observables. Specifically, we prove the following proposition.

Proposition 3 (informal)**.** *(i) Any set of examples that can be correctly classified using a low-rank observable can also be correctly classified using a high-rank observable.*

(ii) There exist sets of examples that can only be correctly classified using an observable of at least a certain rank.

Secondly, we show that by varying the rank of the observables from 1 up to 2^n, we interpolate between the expressivity of linear classifiers on $\mathbb{R}^{2^n}$ up to the expressivity of linear classifiers on $\mathbb{R}^{4^n}$. Specifically, we prove the following proposition.

Proposition 4 (informal)**.** *Let $\mathbb{O}_r$ denote the set of n-qubit observables of rank $\leq r$. Then, we have the following sequence of inclusions on the expressivity of classifier families*

$$\text{linear classifiers on } \mathbb{R}^{2^n} \subseteq \mathcal{C}(\mathbb{O}_1) \subsetneq \mathcal{C}(\mathbb{O}_2) \subsetneq \cdots \subsetneq \mathcal{C}(\mathbb{O}_{2^n}) \subseteq \text{linear classifiers on } \mathbb{R}^{4^n}.$$

Finally, we show that quantum models that use observables with large Frobenius norms can achieve strictly larger margins (i.e., empirical quantities measured on a set of training examples that influence certain generalization bounds) compared to quantum models that use observables with small Frobenius norms. Specifically, we prove the following proposition.

Proposition 5 (informal)**.** *There exist a set of m examples that can only be correctly classified with margin γ using observables of at least Frobenius norm $\gamma\sqrt{m}$.*

In conclusion, by connecting our results to standard structural risk minimization theory, we provide new options for optimally tuning QML models. In particular, our results theoretically motivate two novel kinds of *regularization*, which is a widely used technique to implement structural risk minimization that adds a term to the loss function which penalizes more complex models. Specifically, we theoretically motivate regularizing the rank (for specially designed ansatze) and Frobenius norm of the observables measured.

[§]Here $\sum$ denotes the sum of vector spaces.

References

[1] Vojtěch Havlíček, Antonio D Córcoles, Kristan Temme, Aram W Harrow, Abhinav Kandala, Jerry M Chow, and Jay M Gambetta. Supervised learning with quantum-enhanced feature spaces. *Nature*, 567, 2019.

[2] Maria Schuld and Nathan Killoran. Quantum machine learning in feature Hilbert spaces. *Physical review letters*, 122, 2019.

Graph neural network initialisation for quantum approximate optimisation

Nishant Jain[1], Brian Coyle[2], Niraj Kumar[2], and Elham Kashefi[2,3]
[1] *Indian Institute of Technology, Roorkee*
[2]*School of Informatics, University of Edinburgh, EH8 9AB Edinburgh, United Kingdom*
[3]*Laboratoire d'Informatique de Paris 6, CNRS, Sorbonne Université, 4 place Jussieu, 75005 Paris, France*
Contact e-mail: brian.coyle@ed.ac.uk

Abstract

Approximate combinatorial optimisation is a promising application area for both near-term, and fault-tolerant quantum computers. The quantum approximate optimisation algorithm (QAOA) is one of the foremost algorithms which is suitable to run on noisy intermediate-scale quantum devices, and has seen promising results. In this work, we address two important ingredients in the QAOA for solving the canonical Max-Cut problem. The first is in the initialisation of the circuit parameters, so called 'warm-starting'. Good warm-starting techniques are important to improve the performance and convergence of the algorithm. To this end, we propose graph neural networks (GNNs) for warm-starting the QAOA. This approach has a number of advantages over previous works, in particular, it significantly speeds up initialisation over graph instances, and is capable of generalisation over both problem instance and graph size.

Extended Abstract

In the last several years, there has been tremendous progress in the search for applications for near-term quantum computers, dubbed *noisy intermediate-scale* quantum (NISQ) technology [1]. Among the forerunners for such use cases are *variational* quantum algorithms (VQAs) which began with the variational quantum eigensolver [2] and the quantum approximate optimisation algorithm (QAOA) [3]. In the wake of these, many new algorithms have been proposed tackling problems in a variety of areas [4]. The primary workhorse in such algorithms is typically the parameterised quantum circuit (PQC).

In this work, we focus on one particular VQA, the QAOA, used for approximate discrete combinatorial optimisation, the canonical example of which is Max-Cut on a graph. Here, one aims to partition graph nodes into two sets which have as many edges connecting them as possible. Discrete optimisation problems such as Max-Cut are hard to solve (specifically NP-Hard) and accurate solutions to such problems take exponential time which is generally not feasible. While it is not believed quantum computers can solve such problems *efficiently*, it is hoped that quantum algorithms such as QAOA may be able to outperform classical algorithms by some benchmark. Here, we specifically focus on the *initialisation* of the algorithm, which has shown to have a huge impact on performance [5, 6]. We propose a method to warm-start the algorithm based on using *graph neural networks* (GNNs), which are powerful neural network architectures suited to dealing with graph-specific data [7]. We demonstrate how using GNNs enables high quality solutions, which are generalisable over both problem instance and graph size. Furthermore, once trained, they provide a much faster means of producing initial starting points when compared against previous work [5] using the celebrated Goemans-Williamson algorithm and continuous relaxations of the original Max-Cut problem. We compare against this approach, and also another recent method based on Trotterised quantum annealing (TQA) of [6].

Results We provide numerical implementation of the above proposed method for GNN initialisation of QAOA. One of the key important quantities relevant for benchmarks is the *approximation ratio*, r. This measures the quality of our Max-Cut solution and given by:

$$r = \frac{\text{Approximate cut value}}{\text{Optimal cut value}} \tag{1}$$

Due to the inherent hardness of the Max-Cut problem, we do not believe it is possible for a polynomial time algorithm to achieve r to be arbitrarily close to 1, which would indicate an ideal solution. Recall that the GW algorithm above is a 0.88-approximation algorithm, which means it is guaranteed to produce a cut with a value of r no less than ≈ 0.88.

Graph neural network versus the GW algorithm For the GNN, we use unsupervised training and an architecture similar to [8]. We benchmark the GNN approach to the GW algorithm directly in Fig. 1a by

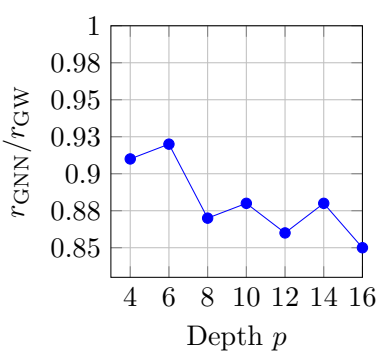

(a) Max-Cut approximation ratios generated by GNN and the GW algorithm, as a function of graph size. r_j is the approximation ratio generated by method $j \in \{\text{GW}, \text{GNN}\}$.

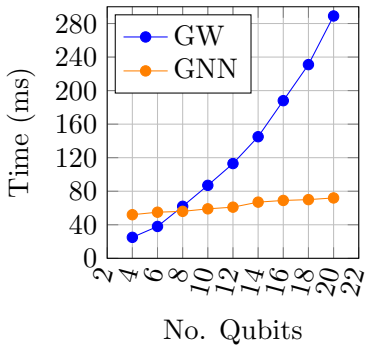

(b) Comparison of time taken by GW algorithm and GNN as a function of qubit number. The GNN provides a large time improvement in Max-Cut inference.

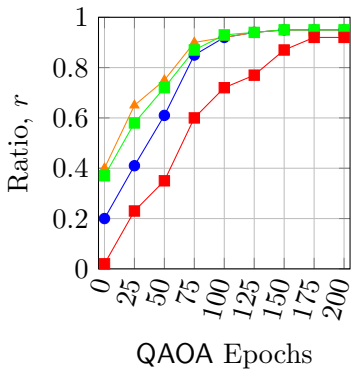

(c) Convergence of initialisation techniques with QAOA on random 8-node 3-regular graphs as a function of training iteration. TQA (blue), Warm-start (orange), GNN (green), Cold-start (red)

Figure 1: Comparison of GNN initilisation to other techniques.

fixing the number of qubits and changing the QAOA depth. Both methods have comparable performance with the GNN approach being able to generate solutions which are between $85 - 90\%$ the quality of those generated by the GW algorithm. However, if then examine Fig. 1b, then the tradeoff becomes apparent. With a small sacrifice in solution quality, the GNN (once trained) is able to generate candidate solutions significantly faster than the GW algorithm, since the GW algorithm must be run separately for each graph instance. The inference time for the GNN is approximately linear, whereas the GW algorithm behaves as some higher degree polynomial (numerically estimated to be $n^{3.5}$ in [9]).

Note that the results have been plotted and the QAOA run for a low qubit number ($n < 25$), which is primarily due to the inherent overhead in the classical simulation of quantum computation. However, it has been numerically verified by [8] that the GNN only continues to improve in quality with increasing graph size (and corresponding qubit number), up to values of 300 nodes or more, even outperforming the GW algorithm at scale. As such, we only expect the GNN approach to improve over the GW algorithm for warm-starting QAOA both in running time, and in solution quality as larger quantum computers become available. We finally compare the GNN initialisation technique against all other techniques in Fig. 1c. We compare against the warm-starting technique using SDP relaxations of [5] ('Warm-start'), and the trotterised quantum annealing ('TQA') based approach of [6]. We also add a random ('Cold-start') initialisation of the parameters as benchmark example.

Train size \ Test size	8	10	12	14
6	0.91	0.93	0.89	0.89
8	0.93			
10		0.93		
12				0.89
14				0.96

Table 1: Value of approximation ratio, r, as a function of training and test graph size. Rows correspond to the graph size on which the model was trained, and the columns correspond to the graph size used for testing.

Generalisation Capabilities of GNNs Here we test the ability of the GNN to generalise in initialising QAOA instances across different problem size. To do so, we train the GNN on small graph instances, and then directly apply it to produce an initialisation for larger instances. We test this for examples of between 6 to 14 nodes in Table 1. Here, we see that training the GNN on small problem instances still is able perform comparably well on large ones. This is a feature which is not available in other initialisation techniques. For the sizes we test, we observe an approximation ratio of at least than 89%, even when only training on a graph half the size of that to be tested.

References

[1] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018.

[2] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5(1):1–7, July 2014.

[3] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A Quantum Approximate Optimization Algorithm. *arXiv:1411.4028 [quant-ph]*, November 2014.

[4] Marco Cerezo, Alexander Poremba, Lukasz Cincio, and Patrick J. Coles. Variational Quantum Fidelity Estimation. *Quantum*, 4:248, March 2020.

[5] Daniel J. Egger, Jakub Mareček, and Stefan Woerner. Warm-starting quantum optimization. *Quantum*, 5:479, June 2021.

[6] Stefan H. Sack and Maksym Serbyn. Quantum annealing initialization of the quantum approximate optimization algorithm. *Quantum*, 5:491, July 2021.

[7] Jie Zhou, Ganqu Cui, Shengding Hu, Zhengyan Zhang, Cheng Yang, Zhiyuan Liu, Lifeng Wang, Changcheng Li, and Maosong Sun. Graph neural networks: A review of methods and applications. *AI Open*, 1:57–81, January 2020.

[8] Weichi Yao, Afonso S. Bandeira, and Soledad Villar. Experimental performance of graph neural networks on random instances of max-cut. In *Wavelets and Sparsity XVIII*, volume 11138, page 111380S. International Society for Optics and Photonics, September 2019.

[9] Martin J. A. Schuetz, J. Kyle Brubaker, and Helmut G. Katzgraber. Combinatorial Optimization with Physics-Inspired Graph Neural Networks. *arXiv:2107.01188 [cond-mat, physics:quant-ph]*, July 2021. arXiv: 2107.01188.

Generation of High Resolution Handwritten Digits with an Ion-Trap Quantum Computer

Manuel S. Rudolph,[1] Ntwali Toussaint Bashige,[2] Amara Katabarwa,[2]
Sonika Johri,[3] Borja Peropadre,[2] and Alejandro Perdomo-Ortiz[1,*]

[1]*Zapata Computing Canada Inc., 325 Front St W, Toronto, ON, M5V 2Y1, Canada*
[2]*Zapata Computing Inc., 100 Federal Street, Boston, MA 02110, USA*
[3]*IonQ Inc., College Park, MD 20740, USA*

Abstract: In this work, we provide the first practical and experimental implementation of a quantum-classical generative algorithm capable of generating high-resolution images of handwritten digits with state-of-the-art gate-based quantum computers. **Pre-print:** arXiv:2012.03924.

In the last decades, machine learning (ML) algorithms have significantly increased in importance and value due to the rapid progress in ML techniques and computational resources [1, 2]. However, even state-of-the-art algorithms face significant challenges in learning and generalizing from an ever increasing volume of unlabeled data [3–5]. With the advent of quantum computing, quantum algorithms for ML arise as natural candidates in the search of applications of noisy intermediate-scale quantum (NISQ) devices, with the potential to surpass classical ML capabilities [6]. Generative ML is one of the most exciting and challenging frontiers in machine learning and among the top contenders for a quantum advantage [7]. Generative models are probabilistic models aiming to capture the most essential features of complex data and to generate similar data by sampling from the trained model distribution. Although quantum generative models have been proven to learn distributions which are outside of classical reach [8–10], it is not clear how NISQ hardware, with their limited number of qubits and level of gate noise, can be best utilized. As such, one of the main challenges that researchers in this field face today is applying and scaling quantum models on small quantum devices to tackle real-world datasets and benchmark possible quantum enhancements. To that end, we introduce the Quantum Circuit Associative Adversarial Network (QC-AAN): a framework combining capabilities of NISQ devices for generative modelling with classical deep learning techniques to learn relevant full-scale data sets (see Figure 1). The framework applies a Quantum Circuit Born Machine (QCBM) [11] to model and re-parameterize the prior distribution of a Generative Adversarial Network (GAN) [12]. This quantum-classical framework exploits the success of GANs in achieving impressive results on high-dimensional datasets, as well as the known dimensionality-reduction capabilities of deep neural networks [7, 13] to implement a quantum generative model as a compact and vital component of the overall algorithm.

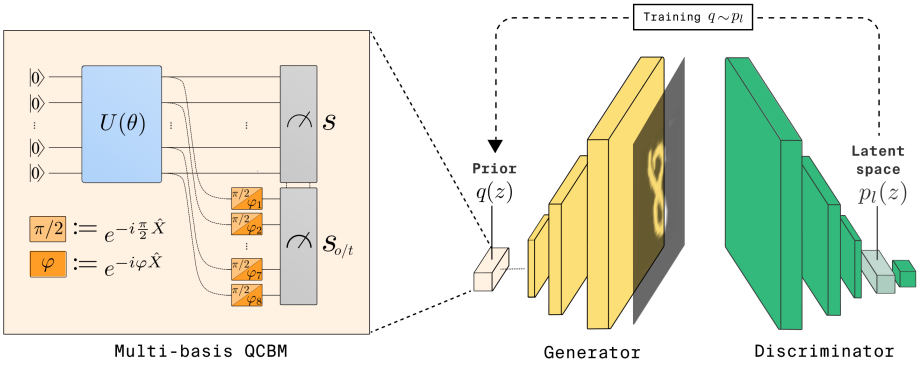

FIG. 1. Schematic description of our QC-AAN framework where the prior of a GAN is modelled by a multi-basis QCBM, which is trained on the latent space distribution during training of the Discriminator.

* alejandro@zapatacomputing.com

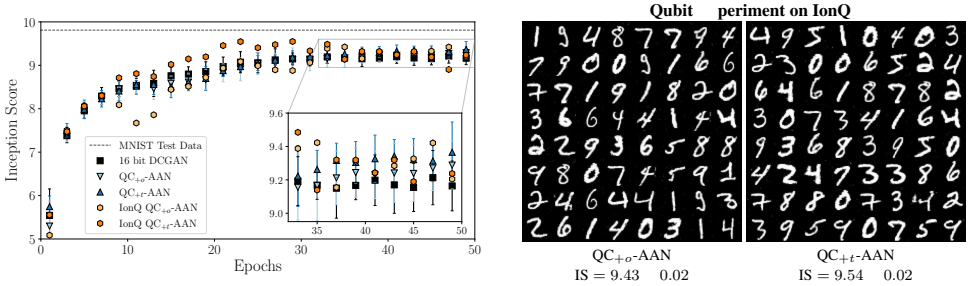

FIG. 2. Left: Quantitative comparison of the QC-AAN against a comparable classical deep convolutional GAN (DCGAN). The experimental realization on the IonQ device includes complete implementation of the multi-basis QCBM on hardware. Right: Images of handwritten digits generated by the experimental implementation of the QC-AAN models on the IonQ device with fixed orthogonal (+o) and trained basis measurements (+t).

The QCBM is a circuit-based generative model which encodes a probability distribution over binary data in the measurement probabilities of a quantum wavefunction. Notably, it can be implemented on most NISQ devices [14–17]. By virtue of being gate-based, QCBM wavefunctions can be measured in multiple bases which can consequently enhance the QC-AAN by providing it with non-classical prior distributions. We call this the multi-basis technique for the QCBM where pre-rotations for quantum measurements can either be trained or fixed to a specific angle, for example to measure in a locally orthogonal basis.
On the classical side, GANs are one of the most popular recent generative machine learning frameworks able to generate remarkably realistic images and other data. In a GAN, a generator G and a discriminator D are trained according to an adversarial loss function with the goal that G learns to generate data which for D are indistinguishable from the real data. The input to G is called the Prior, which is conventionally a continuous uniform or normal distribution with zero mean, although discrete Bernoulli priors have also been shown empirically to be competitive [18]. The AAN framework [19] aims at learning an informed prior distribution which actively takes part in the training of the GAN and can help avoiding common pitfalls of the delicately balanced adversarial game, such as mode collapse and non-convergence [12, 18]. In the QC-AAN, the Prior is modelled by a QCBM which is dynamically trained on the so-called latent space of the Discriminator: a deep and low-dimensional layer reflecting a strongly condensed representation of the data.

We numerically simulate the performance of the QC-AAN with an 8 qubit QCBM using the multi-basis technique and demonstrate an enhancement on the MNIST dataset of handwritten digits, as compared to a comparable classical GAN with strictly the same network architecture. This emphasizes that the choice of a good prior for a deep generative learning task, a topic which is starkly under-studied in the classical ML community, can lead to significant improvements with as few as 8 qubits. Finally, to demonstrate the readiness of our framework, we train the QC-AAN with an experimental implementation of 8 qubits on IonQ's ion-trap quantum computer. To the best of our knowledge, Figure 2 depicts the first practical implementation of a quantum-classical algorithm capable of generating high-resolution digits on a NISQ device. Unlike many other use-case implementations of quantum algorithms on NISQ devices, our models do not under-perform compared to noise-free simulations. This indicates that significant re-parametrization of the prior space, paired with a modest noise floor, provide GANs with an improved trade-off between exploration of the target space and convergence to high-quality data. The QC-AAN framework extends flexibly to more complex data sets, such as data with higher resolution and color, and thus enables a trajectory towards understanding and realizing potential enhancements of state-of-the-art machine learning algorithms by quantum computers.

[1] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, "Deep learning," Nature **521**, 436 – 444 (2015).

[2] Jürgen Schmidhuber, "Deep learning in neural networks: An overview," Neural Networks **61**, 85 – 117 (2015).

[3] Y. Cheng, D. Wang, P. Zhou, and T. Zhang, "Model compression and acceleration for deep neural networks: The principles, progress, and challenges," IEEE Signal Processing Magazine **35**, 126–136 (2018).

[4] Roman Novak, Yasaman Bahri, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein, "Sensitivity and generalization in neural networks: an empirical study," (2018), arXiv:1802.08760 [stat.ML].

[5] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro, "The role of over-parametrization in generalization of neural networks," in *International Conference on Learning Representations* (2019).

[6] John Preskill, "Quantum computing in the NISQ era and beyond," Quantum **2**, 79 (2018).

[7] Alejandro Perdomo-Ortiz, Marcello Benedetti, John Realpe-Gómez, and Rupak Biswas, "Opportunities and challenges for quantum-assisted machine learning in near-term quantum computers," Quantum Science and Technology **3**, 030502 (2018).

[8] Yuxuan Du, Min-Hsiu Hsieh, Tongliang Liu, and Dacheng Tao, "The expressive power of parameterized quantum circuits," arXiv preprint arXiv:1810.11922 (2018).

[9] Ivan Glasser, Ryan Sweke, Nicola Pancotti, Jens Eisert, and J. Ignacio Cirac, "Expressive power of tensor-network factorizations for probabilistic modeling, with applications from hidden Markov models to quantum machine learning," arXiv:1907.03741v2 (2019).

[10] Ryan Sweke, Jean-Pierre Seifert, Dominik Hangleiter, and Jens Eisert, "On the quantum versus classical learnability of discrete distributions," (2020), arXiv:2007.14451 [quant-ph].

[11] Marcello Benedetti, Delfina Garcia-Pintos, Oscar Perdomo, Vicente Leyton-Ortega, Yunseong Nam, and Alejandro Perdomo-Ortiz, "A generative modeling approach for benchmarking and training shallow quantum circuits," npj Quantum Information **5**, 45 (2019).

[12] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, edited by Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger (Curran Associates, Inc., 2014) pp. 2672–2680.

[13] Geoffrey E Hinton and Ruslan R Salakhutdinov, "Reducing the dimensionality of data with neural networks," Science **313**, 504–507 (2006).

[14] Kathleen E. Hamilton, Eugene F. Dumitrescu, and Raphael C. Pooser, "Generative model benchmarks for superconducting qubits," arXiv preprint arXiv:1811.09905 (2018).

[15] Vicente Leyton-Ortega, Alejandro Perdomo-Ortiz, and Oscar Perdomo, "Robust implementation of generative modeling with parametrized quantum circuits," arXiv preprint arXiv:1901.08047 (2019).

[16] Brian Coyle, Maxwell Henderson, Justin Chan Jin Le, Niraj Kumar, Marco Paini, and Elham Kashefi, "Quantum versus classical generative modelling in finance," (2020), arXiv:2008.00691 [quant-ph].

[17] D. Zhu, N. M. Linke, M. Benedetti, K. A. Landsman, N. H. Nguyen, C. H. Alderete, A. Perdomo-Ortiz, N. Korda, A. Garfoot, C. Brecque, L. Egan, O. Perdomo, and C. Monroe, "Training of quantum circuits on a hybrid quantum computer," Science Advances **5** (2019).

[18] Andrew Brock, Jeff Donahue, and Karen Simonyan, "Large scale GAN training for high fidelity natural image synthesis," (2019), arXiv:1809.11096 [cs.LG].

[19] Tarik Arici and Asli Celikyilmaz, "Associative adversarial networks," CoRR (2016), arXiv:1611.06953.

Learning of Quantum PUFs based on single-qubit gates

Anna Pappa[1,2], Niklas Pirnay[1*], and Jean-Pierre Seifert[1]

[1]*Electrical Engineering and Computer Science Department, Technische Universität Berlin, 10587 Berlin, Germany*
[2]*Fraunhofer Institute for Open Communication Systems, Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany*

Abstract

Quantum Physical Unclonable Functions (QPUFs) have been proposed as a way to identify and authenticate quantum devices. In this work we examine classical readout QPUFs based on single qubit rotation gates and show that they are not secure — by demonstrating a Machine Learning based attack on a commercial quantum device.

Secure attestation of cloud computing resources has been in the focus of research to create trust in the cloud, since through it, cloud computing customers are able to ensure that they are provided the correct and trusted underlying hardware [1]. Trusted hardware is especially important for quantum computing, since the high sensitivity to noise in low-grade machines can have detrimental effects on vital calculations. It is therefore an important endeavour to ensure that quantum computing cloud customers can authenticate their quantum devices, to lower the risk of making extensive business decisions based on corrupted results.
In the classical world, Physical Unclonable Functions (PUFs) have been proposed as a way to identify and authenticate electronic devices [2, 3, 4]. Recently, Quantum PUFs (QPUFs) have emerged that aim to achieve the same goals for quantum devices. The primer of Skoric [5] introduced the concept of Quantum Readout PUFs (QR-PUFs). Shortly after, QR-PUFs were generalized and formalized in the framework of Doosti et al. [6] and rigorously analyzed by Arapinis et al. [7]. While these protocols require a QRAM and a quantum channel between the verifier and prover to exchange quantum states, the recent additions by Phalak et al. [8] use only classical communication to provide a secure fingerprint for quantum devices.
In this work, we adapt the QPUF framework of [6] in order to define the category of Classical Readout Quantum PUFs (CR-QPUFs) where verifier and prover communicate classically to authenticate a quantum device. Let us consider the challenge and response scheme $\vec{r}_{\mathrm{out}} = \mathbf{E}_{\mathrm{id}}(U_{\mathrm{in}}(\vec{\theta}))$, where $\vec{r}_{\mathrm{out}}$ is the response, $\mathbf{E}_{\mathrm{id}}$ is the empirical average of the qubits evaluated on the quantum computer with identity id and $U_{\mathrm{in}}(\vec{\theta})$ is the parameterized challenge unitary. Essentially, the expectation value with finite samples $\mathbf{E}_{\mathrm{id}}$ depends on device-specific properties, most notably noise such as gate errors, decoherence and crosstalk. The goal of the CR-QPUF is then to leverage this systematic noise and imperfections to create an unforgeable fingerprint of the quantum device.

We proceed with examining the Hadamard CR-QPUF from [8], where a challenge is defined as

$$U_{\mathrm{in}}(\vec{\theta}) := \bigotimes_{i=1}^{n} H R_Y(\theta_i).$$

We identify essential security flaws in the Hadamard CR-QPUF, which, as we argue, stem from the absence of entanglement, therefore allowing potential attackers to model the QPUF behaviour by learning

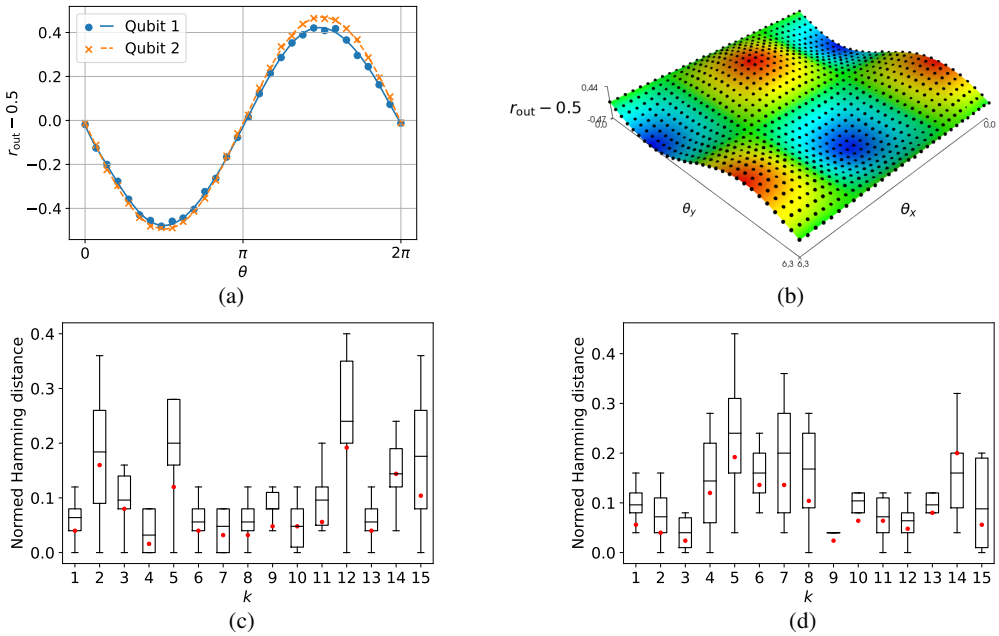

Figure 1 – (a) Measured responses r_{out} and learned model functions for two qubits on `ibmq_belem`. The points are used to learn the low degree polynomials (lines) and predict unknown challenges. The qubits exhibit different characteristics depending on θ, however their expectation value can be accurately learned. (b) Multi-dimensional extension of the Hadamard CR-QPUF for chains of X and Y rotations. The measured responses (black dots) are used to learn the 2D polynomial (surface). (c, d) Attack results for the base and multi-dimensional case respectively. The boxplots show the Hamming distance (Hd) of 15 response signatures in a holdout challenge-response database and their mean. The red dots show the average Hd of the respective predicted responses (obtained through the model functions). Predicted responses get accepted if the Hd is less than or equal to the mean intra Hd in the holdout database.

the characteristics of the individual qubits. We implement the Hadamard CR-QPUF on the publicly available 5-qubit `ibmq_belem` superconducting chip [9] and show that responses can be predicted by learning bounded-degree polynomials [10, 11]. We therefore numerically demonstrate the insufficient security of the Hadamard CR-QPUF as well as natural extensions thereof, where multiple sequential single-qubit rotations are performed. Specifically, in Figure 1 we show the measured responses, learned models and results of the authentication protocol for the one dimensional case and even the multidimensional extension.

Our contributions are concluded with an in-depth discussion of our findings and future CR-QPUF design considerations. In the worst case, to learn the density matrix of an unknown entangled quantum state, an exponential number of measurements are required. However, since the Hadamard CR-QPUF does not use entanglement, the QPUF security reduces to that of one single qubit. This means that an attacker only needs to simulate or predict the behaviour of one qubit at a time. In addition, since the challenges $U_{\text{in}}(\vec{\theta})$ are fixed in their structure, the challenge space is effectively reduced to the rotations $\vec{\theta}$. Hence, one is possibly missing out on leveraging that the concrete structure of random challenges is unknown to an attacker. We believe that these design shortcomings lead to the empirically shown insecurity of the Hadamard CR-QPUF. While the systematic noise present in quantum computers could potentially be used to create an unforgeable fingerprint of the devices, the contradicting incentives, from the one side of quantum computer manufacturers, who want to eliminate systematic noise, and on the other side of end users, who want to use the systematic noise to identify the quantum devices, might hinder the application of QPUF schemes to industrial cloud providers.

References

[1] Dengguo Feng. *Trusted Computing: Principles and Applications*. De Gruyter, 2017.

[2] Ravikanth Pappu, Ben Recht, Jason Taylor, and Neil Gershenfeld. Physical one-way functions. *Science*, 297(5589):2026–2030, 2002.

[3] Christina Brzuska, Marc Fischlin, Heike Schröder, and Stefan Katzenbeisser. Physically uncloneable functions in the universal composition framework. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011*, pages 51–70, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[4] Roel Maes. *Physically Unclonable Functions - Constructions, Properties and Applications*. Springer, 2013.

[5] Boris Škorić. Quantum readout of physical unclonable functions. *International Journal of Quantum Information*, 10(01):1250001, 2012.

[6] Mina Doosti, Niraj Kumar, Mahshid Delavar, and Elham Kashefi. Client-server identification protocols with quantum puf. *arXiv preprint arXiv:2006.04522*, 2020.

[7] Myrto Arapinis, Mahshid Delavar, Mina Doosti, and Elham Kashefi. Quantum physical unclonable functions: Possibilities and impossibilities. *Quantum*, 5:475, 2021.

[8] Koustubh Phalak, Abdullah Ash-Saki, Mahabubul Alam, Rasit Onur Topaloglu, and Swaroop Ghosh. Quantum PUF for Security and Trust in Quantum Computing. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, pages 1–1, 2021. arXiv: 2104.06244.

[9] IBM Quantum. `https://quantum-computing.ibm.com/`. Accessed: 2010-09-30.

[10] Itai Benjamini, Gil Kalai, and Oded Schramm. Noise sensitivity of boolean functions and applications to percolation. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 90(1):5–43, 1999.

[11] Fatemeh Ganji, Domenic Forte, and Jean-Pierre Seifert. Pufmeter a property testing tool for assessing the robustness of physically unclonable functions to machine learning attacks. *IEEE Access*, 7:122513–122521, 2019.