Predicting Features of Quantum Systems using Classical Shadows

Hsin-Yuan Huang
Institute for Quantum Information and Matter,
Department of Computing and Mathematical Sciences
California Institute of Technology, USA
hsinyuan@caltech.edu

Richard Kueng
Institute for Quantum Information and Matter,
Department of Computing and Mathematical Sciences
California Institute of Technology, USA
rkueng@caltech.edu

Abstract

Predicting features of complex, large-scale quantum systems is essential to the characterization and engineering of quantum architectures. The state space of quantum systems is enormous and an exponential number of quantum measurements are required to learn a full description. To overcome this exponential bottleneck, we present an efficient approach for learning a classical model, called the *classical shadow*, of a quantum system from $\log(M)$ measurements only. This number is completely independent of the ambient dimension. Classical shadows can later be used to accurately predict any collection of M linear features. We equip this prediction procedure with rigorous convergence guarantees that saturate fundamental lower bounds from information theory. Numerical experiments support our theoretical findings over a wide range of problem sizes (quantum state space dimensions ranging from 2^2 to 2^{162}) and highlight advantages compared to existing machine learning approaches.

1 Introduction

Learning and characterizing large quantum systems is crucial for the advancement of quantum technologies. The potential applicability of such technologies ranges from quantum sensors based on entangled particles [4, 21] that may improve the detection of gravitational waves to quantum computing devices that solve hard computational problems [18] and may even simulate quantum field theories [10, 15]. Recent, unprecedented advances in the size of controllable quantum systems [14] have pushed traditional techniques for characterizing quantum systems to the limit of their capabilities. These techniques learn a complete description of the quantum system – a procedure known as quantum state tomography [13, 8]. Because an n-qubit quantum system lives in a Hilbert space with dimension 2^n (curse of dimensionality) and quantum measurements are destructive (wavefunction collapse), an exponential number of measurement repetitions is required to fully characterize the system. This is already prohibitive for systems with only a few tens of qubits.

New approaches are required to overcome this fundamental exponential bottleneck. A recent line of research addresses this problem by adapting modern machine learning techniques. The main idea is to perform simple quantum measurements on independent copies of the quantum system in question and train a generative model, such as an auto-regressive recurrent neural network, based on the

Second Workshop on Machine Learning and the Physical Sciences (NeurIPS 2019), Vancouver, Canada.

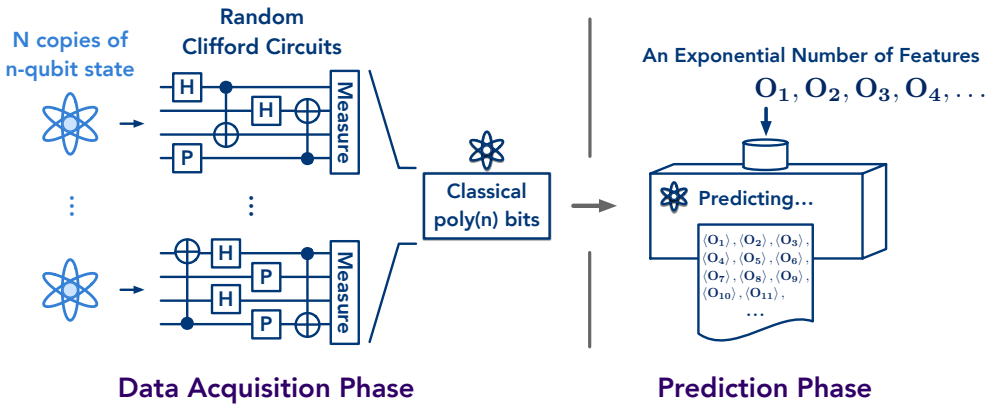

Figure 1: *Caricature of classical shadows:* In the data acquisition phase, we perform random Clifford measurements on independent copies of a n-qubit system. The resulting observations serve as a classical representation of the quantum system – the *classical shadow*. Classical shadows facilitate accurate prediction of a large number of linear features using a simple median-of-means protocol.

observed outcomes [19, 3]. This generative model is then used as an approximate (classical) model of the actual quantum system. Direct access to this model then allows for predicting relevant features. This approach emphasizes tractability and may be potentially useful in characterizing a large class of quantum states using near-term tools. Refs. [19, 3] have supported this approach with empirical evidence. A generative model learned from only a linear number of measurement repetitions (rather than exponential) is capable of accurately predicting many relevant features. In quantum mechanics, interesting features are typically linear functions in the underlying density matrix[1] ρ:

$$o_i(\rho) = \mathrm{trace}(O_i\rho) \quad 1 \leq i \leq M. \tag{1}$$

The fidelity with a pure target state, entanglement witnesses, potential future measurement statistics and expectation values of physical observables are but a few prominent examples. However, like many machine learning approaches, these existing methods are not supported by rigorous guarantees and concrete predictions may not be trustworthy (this will actually be illustrated in our numerical experiments). A more rigorous line of work by Aaronson et al. [1, 2] proves that one can accurately predict many linear quantum features from very few independent copies of the underling quantum system. This approach, however, is extremely (quantum) hardware demanding and requires direct access to an extensive quantum memory.

In this work, we strive for a synthesis of these existing, yet complementary, approaches. In particular, we address the following question: *can we efficiently learn a classical representation of a quantum state from few tractable measurements, such that the representation can later be used to predict quantum features with a rigorous guarantee?* Focusing on data obtained from independent quantum measurement repetitions, we propose an optimal learning procedure that guarantees accurate prediction based on very few state copies and measurements.

2 Main Results

Theorem 1. *There exists a procedure that guarantees the following. Given $B, \epsilon > 0$, The procedure learns a classical representation of an unknown quantum state ρ from independent measurements on $N = \mathcal{O}(B\log(M)/\epsilon^2)$ copies of ρ. Subsequently, given any matrix collection $O_1, \ldots, O_M$ with $B \geq \max_i \mathrm{tr}(O_i^2)$, the procedure can use the classical representation to predict each $\mathrm{tr}(O_i\rho)$ up to additive error ϵ (with high probability).*

Our proof strategy combines the mindset of machine learning (learning an approximate model that provides accurate prediction, not the exact underlying model) with recent insights from traditional quantum state tomography [7] (rigorous convergence guarantees) and the stabilizer formalism [6]

[1]The quantum state of an n-qubit quantum system is fully characterized by a positive semidefinite matrix ρ in $\mathbb{C}^{2^n \times 2^n}$ with unit trace. This object is called the density matrix/operator.

(efficient implementation). This synthesis allows us to rigorously analyze the prediction behavior of an approximate model constructed from very few data points.

The actual procedure is simple and constructive, see Figure 1. In the data acquisition phase, we apply a random quantum circuit, consisting of at most $\mathcal{O}(n^2/\log(n))$ Clifford gates, to each copy of the n-qubit quantum state ρ. Subsequently, we perform a computational basis measurement. Suppose that we perform N repetitions and let U_i denote the i-th quantum circuit, while $\hat{b}_i \in \{0,1\}^n$ is the corresponding measurement outcome. Then the classical representation of the unknown quantum state ρ is simply

$$S(\rho; N) = \{U_1^\dagger |\hat{b}_1\rangle, \ldots, U_N^\dagger |\hat{b}_N\rangle\}$$

which can be stored efficiently using the stabilizer formalism [6]. We call this the *classical shadow* of the quantum state ρ. To predict a certain linear feature $\mathrm{tr}(O\rho)$, we design a predictor that takes into account all the data contained in $S(\rho; N)$. This is reminiscent of nonparametric machine learning algorithms, where the model complexity also scales with the size of the data. Choose $K = \mathcal{O}(\log(M))$, and construct K linear estimators of the underlying density matrix ρ:

$$\hat{\rho}_{(k)} = \frac{2^n+1}{\lfloor N/K \rfloor} \sum_{i=(k-1)\lfloor N/K \rfloor + 1}^{k\lfloor N/K \rfloor} U_i^\dagger |\hat{b}_i\rangle\langle\hat{b}_i| U_i - \mathbb{I} \quad \text{for} \quad 1 \leq k \leq K.$$

Subsequently, we predict a feature $\mathrm{tr}(O\rho)$ by computing $\hat{o} = \mathrm{median}\left\{\mathrm{tr}\left(O\hat{\rho}_{(1)}\right), \ldots, \mathrm{tr}\left(O\hat{\rho}_{(K)}\right)\right\}$. Rich geometric properties of the underlying measurement strategy [20, 22, 11] ensure that this predictor is unbiased with variance bounded by $\mathrm{tr}(O^2)$. Median of means [12, 9] then asserts rapid convergence to the true expectation value.

A natural question is whether our procedure can be further improved. This turns out to be impossible. Fundamental lower bounds from information theory limit the performance of *any* learning procedure that is based on independent quantum measurement repetitions. In other words: linear feature prediction with classical shadows is asymptotically optimal.

Theorem 2. *Any procedure based on independent measurement repetitions requires at least $\Omega(B\log(M)/\epsilon^2)$ copies to guarantee the following. The procedure learns a classical representation from copies of an unknown quantum state ρ. Subsequently, given any M linear features $O_1, \ldots, O_M$ with $B \geq \max_i \mathrm{tr}(O_i^2)$, the procedure can use the classical representation to predict $\mathrm{tr}(O_i\rho)$ up to ϵ-error with high probability.*

3 Numerical Experiments

We have designed classical shadows with tractability in mind. This has allowed us to conduct numerical simulations with more than 160 qubits, i.e. Hilbert space dimension $2^{160} \simeq 1.46 \times 10^{48}$. We compare the performance of classical shadows to existing machine learning approaches. The most recent version of *neural network quantum state tomography* (NNQST) is a generative model that is based on an auto-regressive recurrent neural network trained on independent quantum measurement outcomes (local SIC/tetrahedral POVMs [16]). We consider the task of learning a classical representations of an unknown quantum state and, subsequently, use this to predict the fidelity with some pure target state.

GHZ states In [3], the viability of NNQST is demonstrated by considering GHZ states on a varying number of qubits n. The left-hand side of Figure 2 confirms the linear scaling of NNQST, and the assertion of Theorem 1: classical shadows of *constant* size suffice to accurately estimate a constant number of linear features, regardless of the actual system size. Subsequently, we have also tested the capability of both approaches to detect potential state preparation errors. More precisely, we consider a scenario where the quantum source introduces a phase error with probability $p \in [0,1]$:

$$\rho_p = (1-p)|\psi^+_{\mathrm{GHZ}}(n)\rangle\langle\psi^+_{\mathrm{GHZ}}(n)| + p|\psi^-_{\mathrm{GHZ}}(n)\rangle\langle\psi^-_{\mathrm{GHZ}}(n)|,$$
$$\text{where} \quad |\psi^\pm_{\mathrm{GHZ}}(n)\rangle = \tfrac{1}{\sqrt{2}}\left(|0\rangle^{\otimes n} \pm |1\rangle^{\otimes n}\right).$$

The right-hand side of Figure 2 highlights that classical shadow prediction accurately tracks the decay in target fidelity as the error parameter increases. NNQST, in contrast, seems to consistently overestimate this target fidelity. While existing machine learning approaches are scalable, the prediction may not be trustworthy.

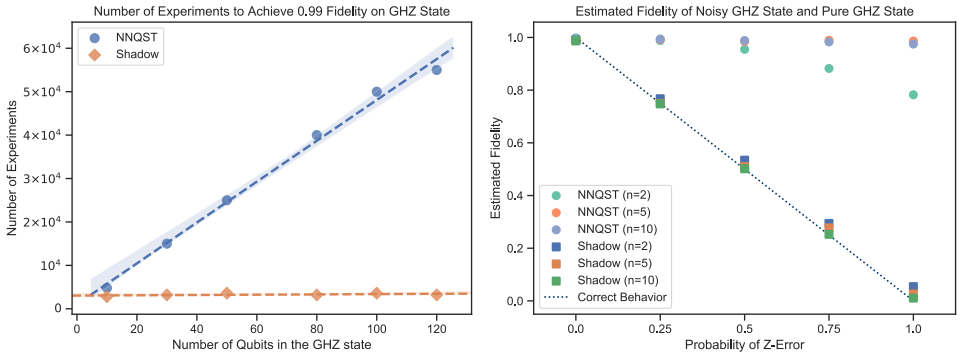

Figure 2: *Comparison between classical shadows and neural network tomography (NNQST); GHZ states.* We learn a classical representation of an unknown quantum state (in this case, a pure or noisy GHZ state). We then test if we have correctly identified this state by predicting the fidelity with a pure GHZ state. NNQST can only estimate classical fidelity (an upper bound on the true quantum fidelity), but we report quantum fidelity for classical shadows.

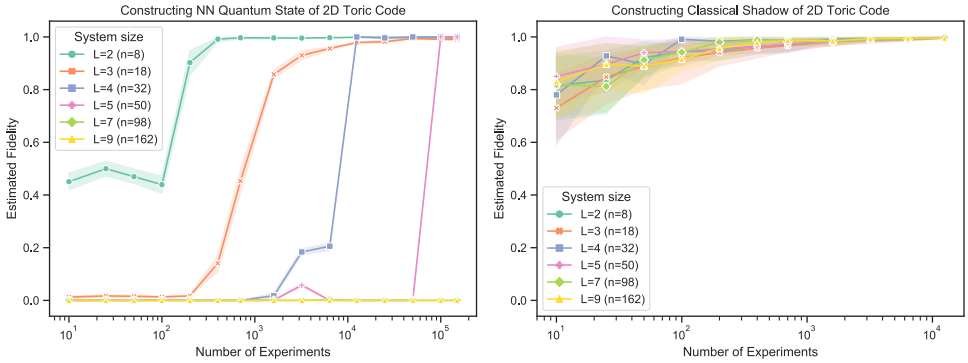

Figure 3: *Comparison between classical shadow and neural network tomography (NNQST); toric code.* We learn a classical representation of an unknown quantum state (in this case, a toric code ground state). We then test if we have correctly identified this state. The shaded regions are the standard deviation of the estimated fidelity over ten runs.

Toric code ground states Toric code ground states are the most prominent example of a topological quantum error-correcting code [5]. The ground state of the toric code [5] is four-fold degenerate and we select the superposition of all closed-loop configurations ($|\psi\rangle \propto \sum_{S:\text{ closed loop}} |S\rangle$) as the target state for both classical shadows and NNQST. The results are shown in Figure 3. Neural network tomography seems to require a number of samples that scales exponentially in the system size n (left). In contrast, the required samples for classical shadows is once more completely independent of the system size. The exponential scaling of NNQST may be rooted in some observed failures of deep learning [17] for learning patterns with combinatorial structures.

4 Outlook

Extension to learning quantum dynamics It would be highly interesting to extend the idea of classical shadows to quantum dynamics and time evolutions. Is it possible to learn an approximate classical model of unknown quantum processes that allow for accurately predicting their action on an exponential number of input states?

Synthesis with modern machine learning An interesting direction would be to combine machine learning and classical shadows, e.g. by training generative models using classical shadows as the input data. A potential benefit comes from the inductive bias imposed by the neural network, which may be biased towards the physical corner of otherwise exponentially large Hilbert spaces. While efficiently

training neural networks with unitary-rotated measurements may be challenging, we believe that a synthesis of machine learning and classical shadows may allow for achieving improved performance with strong theoretical guarantees.

Acknowledgments:

The authors want to thank Victor Albert, Fernando Brandão, John Preskill, Ingo Roth, Joel Tropp, and Thomas Vidick for valuable inputs and inspiring discussions. Leandro Aolita and Giuseppe Carleo provided helpful advice regarding presentation. Our gratitude extends, in particular, to Joseph Iverson who helped us devising a numerical sampling strategy for toric code ground states. HH is supported by the Kortschak Scholars Program. RK acknowledges funding provided by the Office of Naval Research (Award N00014-17-1-2146) and the Army Research Office (Award W911NF121054).

References

[1] S. Aaronson. Shadow tomography of quantum states. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2018, pages 325–338, New York, NY, USA, 2018. ACM.

[2] S. Aaronson and G. N. Rothblum. Gentle measurement of quantum states and differential privacy. In *To appear on the 51th Annual ACM SIGACT Symposium on Theory of Computing*. ACM, 2019.

[3] J. Carrasquilla, G. Torlai, R. G. Melko, and L. Aolita. Reconstructing quantum states with generative models. *Nat. Mach. Intell.*, 1(3):155, 2019.

[4] C. L. Degen, F. Reinhard, and P. Cappellaro. Quantum sensing. *Reviews of modern physics*, 89(3):035002, 2017.

[5] E. Dennis, A. Kitaev, and J. Preskill. Topological quantum memory. volume 43, pages 4452–4505. 2002. Quantum information theory.

[6] D. Gottesman. *Stabilizer codes and quantum error correction. Caltech Ph. D*. PhD thesis, Thesis, eprint: quant-ph/9705052, 1997.

[7] M. Guta, J. Kahn, R. Kueng, and J. A. Tropp. Fast state tomography with optimal error bounds. *arXiv preprint arXiv:1809.11162*, 2018.

[8] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu. Sample-optimal tomography of quantum states. *IEEE T. Inform. Theory*, 63(9):5628–5641, 2017.

[9] M. R. Jerrum, L. G. Valiant, and V. V. Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoret. Comput. Sci.*, 43(2-3):169–188, 1986.

[10] S. P. Jordan, K. S. Lee, and J. Preskill. Quantum algorithms for quantum field theories. *Science*, 336(6085):1130–1133, 2012.

[11] R. Kueng and D. Gross. Qubit stabilizer states are complex projective 3-designs. *arXiv preprint arXiv:1510.02767*, 2015.

[12] A. S. Nemirovsky and D. B. a. Yudin. *Problem complexity and method efficiency in optimization*. A Wiley-Interscience Publication. John Wiley & Sons, Inc., New York, 1983. Translated from the Russian and with a preface by E. R. Dawson, Wiley-Interscience Series in Discrete Mathematics.

[13] R. O'Donnell and J. Wright. Efficient quantum tomography. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 899–912. ACM, 2016.

[14] J. Preskill. Quantum computing in the nisq era and beyond. *Quantum*, 2:79, 2018.

[15] J. Preskill. Simulating quantum field theory with a quantum computer. *arXiv preprint arXiv:1811.10085*, 2018.

[16] J. M. Renes, R. Blume-Kohout, A. J. Scott, and C. M. Caves. Symmetric informationally complete quantum measurements. *J. Math. Phys.*, 45(6):2171–2180, 2004.

[17] S. Shalev-Shwartz, O. Shamir, and S. Shammah. Failures of gradient-based deep learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3067–3075. JMLR. org, 2017.

[18] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.

[19] G. Torlai, G. Mazzola, J. Carrasquilla, M. Troyer, R. Melko, and G. Carleo. Neural-network quantum state tomography. *Nat. Phys.*, 14(5):447, 2018.

[20] Z. Webb. The clifford group forms a unitary 3-design. *arXiv preprint arXiv:1510.02769*, 2015.

[21] S. Zhou, M. Zhang, J. Preskill, and L. Jiang. Achieving the heisenberg limit in quantum metrology using quantum error correction. *Nature communications*, 9(1):78, 2018.

[22] H. Zhu. Multiqubit clifford groups are unitary 3-designs. *Phys. Rev. A*, 96:062336, Dec 2017.