Harvard Data Science Review • Issue 4.1, Winter 2022

Some Recent Progress in Learning Theory: The Quantum Side

Anurag Anshu¹

¹University of California Berkeley, Berkeley, California, United States of America

Published on: Jan 27, 2022

DOI: <https://doi.org/10.1162/99608f92.9c47e402>

License: [Creative Commons Attribution 4.0 International License \(CC-BY 4.0\)](https://creativecommons.org/licenses/by/4.0/)

Keywords: quantum machine learning, quantum learning theory quantum property testing, quantum spin-systems

1. Introduction

The review (Wang, 2022) covers the exciting area of quantum machine learning, where major efforts are underway in finding examples of machine learning problems where quantum algorithms can give speedups over classical algorithms (termed as ‘quantum advantage’). However, there are several challenges in this direction. For low-rank linear algebra problem instances, a series of results initiated by Tang (2019) have shown barriers to a clear quantum advantage. Training quantum neural networks and related objects have seen the problem of barren plateau while performing gradient descent (McClean et al., 2018). Rigorous quantum advantages are beginning to emerge (Liu et al., 2021), but there are several milestones to cover before the same in practical questions can be found (Aaronson, 2015). Probably approximately correct (PAC) learning of concept classes has been studied through the quantum lens, with results (Arunachalam & de Wolf, 2017; Arunachalam et. al., 2020) showing limitations on the reach of quantum algorithms over the classical.

Another line of work that aims to learn properties of unknown quantum systems in an efficient manner has seen many rigorous results. Here, the race is not with classical computers, rather with the quantum formalism itself: can we avoid performing full tomography of an unknown quantum state (which on n qubits would require exponential in n many samples from the unknown state) if we are interested in gaining a specific information about the state? Readers familiar with (classical) learning theory can place this question alongside some well-known theoretical frameworks, such as property testing and statistical inference in graphical models. Quantum mechanics offers new challenges, hiding relevant information in the vastness of the Hilbert space or in entanglement across qubits, which cannot be captured by known classical techniques. Still, important lessons from learning theory are all well embraced, as the reader may note in the discussions below.

2. Results

A prototypical task in this regime is shadow tomography of Aaronson (2020). A quantum state provides a full description of a quantum system, specifying the measurement statistics on any quantum measurement. But converging onto a pretty accurate description of a quantum system, on which we have no prior, is very expensive, as hinted earlier. The goal of shadow tomography is to obtain a pretty good description with respect to just a few specified quantum measurements; a goal relevant to the practical and computationally bounded scenarios. Aaronson shows that there is a scheme to output a quantum state that requires polynomially many (in the number of qubits) samples, as long as the number of specified measurements are at most exponential in number. While the doubly exponential number of ‘almost all’ measurements¹ are left uncertain, one hardly ever explores a Hilbert space in such details.

Aaronson’s shadow tomography scheme is not time efficient, a question that is addressed in the work of Huang et al. (2020). The authors perform efficient quantum operations on the quantum samples and construct a description of the quantum state that gives an accurate estimate for several quantum measurements (especially measurements that have low rank). The guarantees, however, depend on the so called ‘shadow norm’ of the quantum measurements and may sometimes be exponentially large. An important question that remains is to find a scheme that is time efficient and achieves polynomial sample complexity on any set of exponentially many measurements.

The above framework is built on the premise that the observer has no prior knowledge of the unknown quantum state. However, physical scenarios involve quantum states that are much more structured; even nature prefers not to explore the Hilbert space in full details. Assuming that the quantum samples are from a ‘physical’ family of quantum state, there is a hope that learning can be achieved efficiently.² Physicists’ favorite place to find such quantum states is in spin systems, a collection of particles that interact locally as dictated by physics. When the particles are arranged on a line forming the so-called ‘spin-chains,’ the canonical physical quantum states are called matrix product states (MPS). A pioneering work (Cramer et al., 2010) showed that MPS can be learned sample and time efficiently, as long as their ‘bond dimension’ is small. The work also constructed a heuristic algorithm that performed very simple quantum measurements involving a few qubits, a desirable feature from the experimental point of view. An important open question here is the learnability of tensor network states, which are the generalizations of MPS on the more complex (but also more prevalent) higher dimensional systems and may be viewed as the quantum analogues of Hidden Markov models.

Another class of quantum states on spin systems are the quantum Gibbs states. A quantum Gibbs state describes the state of a system that is in contact with a heat bath. The mathematical expression for Gibbs state nicely encodes the interaction between the particles and is written as an exponential of the Hamiltonian (the matrix describing the interactions). It generalizes the graphical models or Ising distributions, which are widely used in machine learning literature (Bresler, 2015; Chow & Liu, 1968; Hinton et al., 1986; Klivans & Meka, 2017). Learning quantum Gibbs states is equivalent to learning their Hamiltonian, for which an algorithm was presented in Anshu et al. (2021). The sample complexity of the algorithm is polynomial in the number of particles and is time efficient whenever the quantum partition function can be efficiently computed. More recently, a time- and sample-optimal algorithm has been presented (Haah et al., 2021) in the high-temperature regime above the phase transition point. Both the algorithms perform simple few-qubit measurements on the quantum samples, making them experimentally feasible. Finding a time-efficient algorithm is wide open in the regime of lower temperatures, which arise in various settings that include the entanglement Hamiltonians (Kokail et al., 2021) and effective Hamiltonians. Yet another natural question is the verification of quantum Gibbs states, which has been studied in the classical setting as an identity testing problem (Daskalakis et al., 2018). It may also find applications in the calibration of quantum Gibbs samplers, which are widely used in quantum algorithms (Brandão & Svore, 2017; Brandão et al., 2019).

The field presents many more exciting challenges that were not detailed above. The learnability or identity testing of quantum states that can be generated by low-depth quantum computations is wide open. They are the focus of research efforts in near-term quantum algorithms (Farhi et al., 2014; Peruzzo et al., 2014) and offer provable advantage over similar classical computations (Bravyi et al., 2018). A series of results have taken the property testing view (Bădescu et al., 2019; Bubeck et al., 2020; O’Donnell & Wright, 2015), offering significant advantage over full tomography (Haah et al., 2017; O’Donnell & Wright, 2015, 2016) in various tasks. Progress on wider family of property testing questions, and especially understanding the role of entangling measurements (Bubeck et al., 2020), remains an important avenue for research. In summary, it is an exciting time to think deeply about the learnability of quantum systems and push the current algorithms to a level of efficiency that can be implemented in practice.

Acknowledgments

I thank Srinivasan Arunachalam, Xun Gao, and Sona Najafi for helpful feedback on the commentary.

Disclosure Statement

Anurag Anshu has no financial or non-financial disclosures to share for this article.

References

- Aaronson, S. (2015). Read the fine print. *Nature Physics*, 11(4), 291–293. <https://doi.org/10.1038/nphys3272>
- Aaronson, S. (2020). Shadow tomography of quantum states. *SIAM Journal on Computing*, 49(5), 368–394. <https://doi.org/10.1137/18M120275X>
- Anshu, A., Arunachalam, S., Kuwahara, T., & Soleimanifar, M. (2021). Sample-efficient learning of interacting quantum systems. *Nature Physics*, 17(8), 931–935. <https://doi.org/10.1038/s41567-021-01232-0>
- Arunachalam, S., & de Wolf, R. (2017). Guest column: A survey of quantum learning theory. *SIGACT News*, 48(2), 41–67. <https://doi.org/10.1145/3106700.3106710>
- Arunachalam, S., Grilo, T. Gur, A. B., Oliveira, I. C., & Sundaram, A. (2020). Quantum learning algorithms imply circuit lower bounds. *arXiv*. <https://doi.org/10.48550/arXiv.2012.01920>
- Bădescu, C., O’Donnell, R., & Wright, J. (2019). Quantum state certification. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (pp. 503–514). Association for Computing Machinery. <https://doi.org/10.1145/3313276.3316344>
- Brandão, F. G. S. L., Kalev, A., Li, T., Lin, C. Y.-Y., Svore, K. M., & Wu, X. (2019). Quantum SDP solvers: Large speed-ups, optimality, and applications to quantum learning. In C. Baier, I. Chatzigiannakis, P. Flocchini,

& S. Leonardi (Eds.), *Leibniz International Proceedings in Informatics (LIPIcs): Volume 132. Proceedings of the 46th International Colloquium on Automata, Languages, and Programming* (Article 14). Dagstuhl Publishing. <https://drops.dagstuhl.de/opus/volltexte/2019/10603/pdf/LIPIcs-ICALP-2019-27.pdf>

Brandão, F. G. S. L., & Svore, K. M. (2017). Quantum speed-ups for solving semidefinite programs. In *Proceedings of the 58th Annual Symposium on Foundations of Computer Science*, (pp. 415–426). IEEE. <https://doi.org/10.1109/FOCS.2017.45>

Bravyi, S., Gosset, D., & Konig, R. (2018). Quantum advantage with shallow circuits. *Science*, 362(6412), 308–311. <https://doi.org/10.1126/science.aar3106>

Bresler, G. (2015). Efficiently learning Ising models on arbitrary graphs. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing* (pp. 771–782). Association for Computing Machinery. <https://doi.org/10.1145/2746539.2746631>

Bubeck, S., Chen, S., & Li, J. (2020). Entanglement is necessary for optimal quantum property testing. In *Proceedings of the 61st Annual Symposium on Foundations of Computer Science* (pp. 692–703). IEEE. <https://doi.org/10.1109/FOCS46700.2020.00070>

Chow, C., & Liu, C. (1968). Approximating discrete probability distributions with dependence trees. *IEEE Transactions on Information Theory*, 14(3), 462–467. <https://doi.org/10.1109/TIT.1968.1054142>

Cramer, M., Plenio, M. B., Flammia, S. T., Somma, R., Gross, D., Bartlett, S. D., Landon-Cardinal, O., Poulin, D., & Liu, Y.-K. (2010). Efficient quantum state tomography. *Nature Communications*, 1(1), Article 149. <https://doi.org/10.1038/ncomms1147>

Daskalakis, C., Dikkala, N., & Kamath, G. (2018). Testing Ising models. *IEEE Transactions on Information Theory*, (65)11, 1989–2007. <https://doi.org/10.1109/TIT.2019.2932255>

Farhi, E., Goldstone, J., & Gutmann, S. (2014). A quantum approximate optimization algorithm. *arXiv*. <https://doi.org/10.48550/arXiv.1411.4028>

Haah, J., Harrow, A. W., Ji, Z. Wu, X., & Yu, N. (2017). Sample-optimal tomography of quantum states. *IEEE Transactions on Information Theory*, 63(9), 5628–5641. <https://doi.org/10.1109/TIT.2017.2719044>

Haah, J., Kothari, R., & Tang, E. (2021). Optimal learning of quantum Hamiltonians from high-temperature Gibbs states. *arXiv*. <https://doi.org/10.48550/arXiv.2108.04842>

Hinton, G. E., & Sejnowski, T. J. (1986) Learning and relearning in Boltzmann machines. In D. E. Rumelhart, & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Volume 1: Foundations* (pp. 282–317). MIT Press.

<https://papers.cnl.salk.edu/PDFs/Learning%20and%20Relearning%20in%20Boltzmann%20Machines%201986-3239.pdf>

Huang, H.-Y., Kueng, R., & Preskill, J. (2020). Predicting many properties of a quantum system from very few measurements. *Nature Physics*, 16(10), 1050–1057. <https://doi.org/10.1038/s41567-020-0932-7>

Klivans, A. R., & Meka, R. (2017). Learning graphical models using multiplicative weights. In *Proceedings of the 58th Annual Symposium on Foundations of Computer Science* (pp. 343–354). IEEE. <https://doi.org/10.1109/FOCS.2017.39>

Kokail, C., van Bijnen, R., Elben, A., Vermersch, B. & Zoller, P. (2021). Entanglement Hamiltonian tomography in quantum simulation. *Nature Physics*, 17(8), 936–942. <https://doi.org/10.1038/s41567-021-01260-w>

Liu, Y., Arunachalam, S., & Temme, K. (2021). A rigorous and robust quantum speed-up in supervised machine learning. *Nature Physics*, 17(9), 1013–1017. <https://doi.org/10.1038/s41567-021-01287-z>

McClean, J. R., Boixo, S., Smelyanskiy, V. N., Babbush, R., & Neven, H. (2018). Barren plateaus in quantum neural network training landscapes. *Nature Communications*, 9(1), Article 4812. <https://doi.org/10.1038/s41467-018-07090-4>

O’Donnell, R., & Wright, J. (2015). Quantum spectrum testing. In *Proceedings of the Forty-Seventh Annual ACM Symposium on Theory of Computing* (pp. 529–538). Association for Computing Machinery. <https://doi.org/10.1145/2746539.2746582>

O’Donnell, R., & Wright, J. (2016). Efficient quantum tomography. In *Proceedings of the Forty-Eighth Annual ACM Symposium on Theory of Computing*, (pp. 899–912). Association for Computing Machinery. <https://doi.org/10.1145/2897518.2897544>

Peruzzo, A., McClean, J., Shadbolt, P., Yung, M.-H., Zhou, X.-Q., Love, P. J., Aspuru-Guzik, A., & O’Brien, J. L. (2014). A variational eigenvalue solver on a photonic quantum processor. *Nature Communications*, 5, Article 4213. <https://doi.org/10.1038/ncomms5213>

Tang, E. (2019). A quantum-inspired classical algorithm for recommendation systems. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing* (pp. 217–228). Association for Computing Machinery. <https://doi.org/10.1145/3313276.3316310>

Wang, Y. (2022). When quantum computation meets data science: Making data science quantum. *Harvard Data Science Review*, 4(1). <https://doi.org/10.1162/99608f92.ef5d8928>

©2022 Anurag Anshu. This article is licensed under a Creative Commons Attribution (CC BY 4.0) [International license](#), except where otherwise indicated with respect to particular material included in the article.

Footnotes

1. This number is similar to the number of subsets of n -bit strings, each of which form an event. [↩](#)
2. We will leave unanswered the question of what makes a family of states physical. A tolerable answer would be anything that physicists encounter in their studies and the examples discussed here are all of this flavor. [↩](#)