Journal Club for Condensed Matter Physics DOI:10.36471/JCCM_May_2022_01
https://www.condmatjclub.org

Learning from memories of quantum states

Quantum advantage in learning from experiments
Authors: Hsin-Yuan Huang, Michael Broughton, Jordan Cotler, Sitan Chen, Jerry Li, Masoud Mohseni, Hartmut Neven, Ryan Babbush, Richard Kueng, John Preskill, Jarrod R. McClean
arXiv:2112.00778

Recommended with a Commentary by Ehud Altman, UC Berkeley

Most experiments with physical quantum systems are based on estimating expectation values of observables in a quantum state. While expectation values can only be estimated by averaging over repeated measurements due to the inherent randomness of quantum measurements, in most experiments this need for averaging does not pose a serious problem. The observables of interest usually have an expectation value of order one, such that after averaging over not too many measurements we can estimate the expectation value to sufficient accuracy. However, a quantum state generally encodes far more information than can be extracted from expectation values of simple observables. For example, to determine if the system generating the state had previously interacted with the environment we may want to estimate the purity $\mathrm{tr}(\rho^2)$, which cannot be represented as an observable expectation value. Measuring the purity to fixed accuracy ϵ generally requires a number of repeated measurements scaling exponentially with system size. Below we give more examples of interesting properties that cannot be estimated in reasonable time by repeated measurements of the same state. This raises the question, whether there are more efficient ways to mine information that is buried deep in the quantum state.

The recommended paper shows that the effort needed to learn certain aspects of a many-body quantum state generated in an experiment can be reduced exponentially by utilizing quantum processing. More specifically, the paper compares two distinct paradigms for extracting information from experiments. The standard (or ”classical”) approach is to repeat the experiment multiple times and obtain a final result through statistical analysis of the measurement data collected after each run. In the alternative (”quantum”) paradigm the experiment is also repeated, but instead of performing a measurement after each run, the entire quantum state of the system is loaded into a quantum memory until there are k identical copies of the state on the register. Measurements are performed only after some quantum processing is applied to the replicated state. Then the entire protocol can be repeated. In the simplest examples to be discussed below, it is enough to have two copies of the state and the quantum processing involved is nothing more than measuring two qubit observables (i.e. correlations) between the two copies. The idea that certain information about a quantum state can be obtained exponentially faster using the second experimental paradigm was developed theoretically in previous papers [1, 2]. The recommended paper extends the new paradigm to a somewhat broader set of questions and presents an experimental proof of principle of the idea.

The essence of the quantum advantage demonstrated in the paper can be elucidated with the following example. Suppose a physical system with n qubits consistently generates the same state, described by the density matrix ρ. Our experimental task is to decide if ρ is a maximally mixed state or a random pure state. In the standard measurement paradigm we measure all n qubits after each run. Since getting any given bit string is exponentially rare in n, detecting a difference between the two paradigms requires more than $2^{n/2}$ repetitions of the experiment [*]. The task becomes exponentially easier if we can load two replicas of the state on a quantum register. In this case we can accomplish the task by measuring the parity of the register state under swap of the two replicas in what is known as the "swap test". If the experiment generated a pair of identical pure states, then the register is invariant to swapping them. We are then guaranteed to measure even parity. On the other hand, if we were handed maximally mixed states, swapping them without swapping the environments that they are entangled with isn't a symmetry. We would measure even and odd parity with equal probability. Note that measuring the total parity is in principle easy because it is the sum of the parity eigenvalues of local observables $\Pi_i = \frac{1}{2}(1 + \boldsymbol{\sigma}_{1,i} \cdot \boldsymbol{\sigma}_{2,i})$. This is a two qubit operator spanning qubits corresponding to different replicas, but encoding the same site i of the physical system.

Huang *et. al.* describe several tasks which have a similar flavor, but that are somewhat less contrived than the above example. One task, demonstrated experimentally using Google's Sycamore processor, involved learning the symmetry class of an unknown physical process that generates an n-qubit state. Specifically, having access to the output state, the task is to decide if it was generated by a time evolution in the unitary class or the orthogonal class. In the standard paradigm one can accomplish this by estimating the expectation value of a purely imaginary operator, such as σ^y, which should be exactly zero if the time evolution was in the orthogonal class. However, distinguishing zero from an expectation value exponentially small in n that would obtain after generic unitary evolution, would seem to require more than 2^n repeated measurements. Indeed the recommended paper proves that no protocol can overcome this exponential barrier without a quantum memory. Remarkably, in the quantum enhanced protocol the symmetry is identified with very few repetitions using a simple variation of the swap test described above.

The final task discussed in the paper, termed quantum principal component analysis (PCA), may be more natural and is not phrased as hypothesis testing like the previous examples. The task can be related to a common situation, where we are interested in the properties of a pure state $|\psi_0\rangle$, yet we only have access to some polluted version of it, represented by the mixed state ρ. A natural example in which $|\psi_0\rangle$ is guaranteed to be the principal component of ρ (i.e. the eigenvector with largest eigenvalue) is when $|\psi_0\rangle$ is the ground state, whereas ρ the corresponding finite temperature state of the system under study. The recommended paper proves that in some cases inferring expectation values in $|\psi_0\rangle$ from measurements made in ρ requires more than $2^{n/2}$ even if the spectrum of ρ has a constant gap below the largest eigenvalue. On the other hand, it was previously shown in [3] that this task can be performed with linear effort if we have access to a quantum register

---

[*] A non-trivial proof for the lower bound $2^{n/2}$ is given in the recommended paper

hosting multiple copies of ρ.

The algorithm proposed in Ref. [3] is more complicated than the previous examples and was not demonstrated experimentally. However a closely related idea of "quantum virtual cooling" was recently demonstrated in an ultra-cold atom system [4]. The basic idea is that using a replicated state $\rho^{\otimes k}$ and the ability to perform operations exchanging the different copies, one can measure expectation values in the state $\rho^k/\mathrm{tr}(\rho^k)$. Thus, operating on a quantum register hosting k copies of a thermal state gives access to correlations at a temperature T/k.

At this point, gaining significant quantum advantage in conventional computational tasks still seems far in the future and inaccessible to near term noisy quantum devices. However, the recommended paper suggests a somewhat different application for quantum processing, which may offer a more immediate advantage: extracting information on the properties of a quantum state and, by extension, learning about the physical system that generated it. The crucial element of the scheme is a quantum memory that can store multiple copies of the state before any measurements are performed. With multiple copies of the state in hand one can perform measurements of observables that span different copies, which would not be possible with access to only a single copy of the state at a time. It is worth noting that the specific tasks discussed in the paper are somewhat unusual and it is not clear to what extent the algorithms proposed to achieve them are scalable in presence of noise. Nevertheless, the fact that a significant advantage could be demonstrated using a noisy processor with 40 physical qubits is a good sign that useful applications are within reach. Finally, the quantum learning protocol described here could have even more real world impact if a similar advantage is established for learning properties of classical distributions pertinent for machine learning tasks.

References

[1] D. Aharonov, J. Cotler, and X.-L. Qi, "Quantum algorithmic measurement," arXiv:2101.04634 (2021).

[2] S. Chen, J. Cotler, H.-Y. Huang, and J. Li, "Exponential separations between learning with and without quantum memory," arXiv:2111.05881

[3] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum principal component analysis," Nat. Phys. 10, 631 (2014).

[4] J. Cotler, S. Choi, A. Lukin, H. Gharibyan, T. Grover, M. E. Tai, M. Rispoli, R. Schittko, P. M. Preiss, A. M. Kaufman, et al., "Quantum virtual cooling," Phys. Rev. X 9, 031013 (2019).