Some Open Problems in Quantum Learning

Sitan Chen, Jordan Cotler, Hsin-Yuan Huang, Jerry Li, Ewin Tang

October 27, 2024

Memory Lower Bounds

Entanglement Testing

NISQ Lower Bounds

Hierarchy for Replica Quantum Advantage

Shadow Tomography

Instance-Optimal State Certification

Lower Bounds for Quantum Channel Property Testing

Learning Hamiltonians from Long Time Evolutions

Lower Bounds for Hamiltonian Learning from Real-Time Evolution

Does Amplitude Amplification Require Inverses?

Polemical Discussion

Memory Lower Bounds

We have an example of a learning problem which requires $\Omega(2^{n-k})$ queries for $n+k$ qubits of quantum memory (Chen-Cotler-Huang-Li, Huang et al., Chen-Gong-Ye). But all existing methods for proving quantum memory lower bounds for statistical tasks fail to say anything when the algorithm has more than $2n$ qubits of quantum memory.

Problem

Are there methods for proving strong lower bounds in this regime? More generally, what kinds of results can we prove when we have constraints on both the amount of classical memory and the amount of quantum memory we have?

Chen, Sitan, Cotler, Jordan, Huang, Hsin-Yuan, & Li, Jerry. "Exponential separations between learning with and without quantum memory." FOCS (2021).

Huang, Hsin-Yuan, Broughton, Michael, Cotler, Jordan, Chen, Sitan, Li, Jerry, et al. "Quantum advantage in learning from experiments." Science 376.6598 (2022): 1182-1186.

Chen, Sitan, Gong, Weiyuan, & Ye, Qi. "Optimal tradeoffs for estimating Pauli observables." arXiv preprint arXiv:2404.19105 (2024).

Entanglement Testing

A bipartite mixed state on $2n$ qubits is said to be *separable* if it can be written as a convex combination of states of the form $\rho \otimes \sigma$, where ρ and σ are n-qubit density matrices.

Problem

Given copies of a state which is promised to be either separable or ϵ-far in trace distance from any separable state, how many copies are needed to determine this?

For entangled measurements, the only known upper bound is $2^{4n}/\epsilon^2$ via full state tomography, and the only known lower bound is $2^{2n}/\epsilon^2$, based on lower bounds for quantum identity testing by Bădescu-O'Donnell '20.

Bădescu, Costin, & O'Donnell, Ryan. "Lower bounds for testing complete positivity and quantum separability." LATIN. Springer, 2020. 375-386.

NISQ Lower Bounds

Chen-Cotler-Huang-Li defined the NISQ complexity class and established (i) NISQ query lower bounds for certain oracle problems (Grover, shadow tomography), and (ii) the construction of an oracle O for which $\mathsf{BPP}^O \subsetneq \mathsf{NISQ}^O \subsetneq \mathsf{BQP}^O$. These results motivate the following two questions:

Problem (NISQ lower bounds)

1. *What is a NISQ query lower bound for Simon's problem? For Forrelation?*
2. *Can one construct a more 'natural' oracle O for which $BPP^O \subsetneq NISQ^O \subsetneq BQP^O$?*

Chen, Sitan, Cotler, Jordan, Huang, Hsin-Yuan, & Li, Jerry. "The complexity of NISQ." Nature Communications 14.1 (2023): 6001.

Hierarchy for Replica Quantum Advantage

Aharonov-Cotler-Qi showed that there is a property testing problem (i.e. distinguishing between a fixed Haar-random pure state and a maximally mixed state for n qubits) which requires exponentially queries if a quantum computer can only measure one sample at a time, but requires only $O(1)$ queries if the quantum computer can make joint measurements on two samples (i.e. *replicas*) at a time. We say that this is a learning task which is 1-*replica hard* and 2-*replica easy*.
In Chen-Cotler-Huang-Li '21, we construct learning problems which are k-*replica hard* and poly(n, k)-*replica easy*. We thus ask the following:

Problem

Are there any learning problems which are k-replica hard but $k+1$ (or $k+O(1)$)-replica easy for any $k>1$?

Aharonov, Dorit, Cotler, Jordan, & Qi, Xiao-Liang. "Quantum algorithmic measurement." Nature Communications 13.1 (2022): 1-9.

Chen, Sitan, Cotler, Jordan, Huang, Hsin-Yuan, & Li, Jerry. "A hierarchy for replica quantum advantage." arXiv:2111.05874 (2021).

Shadow Tomography

Shadow tomography is the following task: given observables (i.e. Hermitian operators) $O_1, \ldots, O_m \in \mathbb{C}^{d\times d}$ and copies of an unknown quantum state $\rho \in \mathbb{C}^{d\times d}$, estimate $\mathrm{Tr}(O_i\rho)$ for $i = 1, \ldots, m$ to within error ϵ. If $O_1, \ldots, O_m$ and ρ were "classical," i.e. if they were diagonal, then this is simply the question of estimating *m statistical queries* of a distribution on d elements, which has sample complexity $\log(m)/\epsilon^2$ by Chernoff and union bound.
What is the optimal sample complexity in the quantum setting? The best known upper bound currently, due to Bădescu-O'Donnell '21 and Bene Watts-Bostanci '22, to, is $O(\log^2(m)\log(d)/\epsilon^4)$.

Problem

Can one improve this to match the classical bound, or is there a provable separation?

Bădescu, Costin, & O'Donnell, Ryan. "Improved quantum data analysis." STOC (2021): 1398-1411.

Watts, Adam Bene, & Bostanci, John. "Quantum event learning and gentle random measurements." arXiv:2210.09155 (2022).

Instance-Optimal State Certification

Quantum state certification is the natural quantum analogue of the classical problem of identity testing. Formally, it asks: given a reference state $\sigma \in \mathbb{C}^{d \times d}$ and copies of an unknown state ρ, determine whether $\rho = \sigma$ or whether $\|\rho - \sigma\|_{\mathsf{tr}} > \epsilon$.
For worst-case σ, it is well-understood what the sample complexity for this task is, and we have tight rates. But in general, one can ask for *instance-optimal* bounds, i.e. sample complexity bounds that depend on the reference state σ. For unentangled measurements, it was shown by Chen-Li-O'Donnell '22 that the sample complexity is essentially $F(\sigma, \mathbb{I}/d) \cdot d^{3/2}/\epsilon^2$, where F is the *fidelity* between quantum states.

Problem

Is there an analogous result for entangled measurements?

Chen, Sitan, Huang, Brice, Li, Jerry, & Liu, Allen. "Tight bounds for quantum state certification with incoherent measurements." arXiv preprint arXiv:2204.07155 (2022).

Chen, Sitan, Li, Jerry, & O'Donnell, Ryan. "Toward instance-optimal state certification with incoherent measurements." arXiv:2102.13098 (2021).

O'Donnell, Ryan, & Wright, John. "Quantum spectrum testing." STOC (2015): 529-538.

Lower Bounds for Quantum Channel Property Testing

The following learning task was studied in Aharonov-Cotler-Qi and Chen-Cotler-Huang-Li. The task is to determine whether an unknown channel is either (i) the maximally depolarizing channel, or (ii) a fixed Haar-random unitary channel. If the quantum computer can only adaptively prepare a quantum state, apply the channel, and then adaptively make a measurement, then $\Omega(2^{n/3})$ queries are required to distinguish between (i) and (ii), and $O(2^{n/2})$ queries are necessary. If the quantum computer can query the channel multiple times to make a joint measurement on the outputs of two channels, then (i) and (ii) can be distinguished with $O(1)$ queries. There is a generalization of this learning task for distinguishing quantum channels belonging to certain symmetry classes. We have the following two questions:

Problem

1. *There is a gap between the $\Omega(2^{n/3})$ query lower bound and the $O(2^{n/2})$ query upper bound. Can one prove a tight bound of $\Theta(2^{n/2})$?*
2. *Are there other natural property testing problems (for which one can prove interesting results) that do not involve Haar-random unitaries (or their psuedo-random proxies)? Some progress in this direction involves learning Pauli channels (see e.g. Chen-Zhou-Seif-Jiang '22). Natural targets include classes of low-depth quantum circuits, or classes of local Hamiltonians.*

Aharonov, Dorit, Cotler, Jordan, & Qi, Xiao-Liang. "Quantum algorithmic measurement." Nature Communications (2022).

Chen, Sitan, Cotler, Jordan, Huang, Hsin-Yuan, & Li, Jerry. "Exponential separations between learning with and without quantum memory." FOCS (2021).

Chen, Senrui, Zhou, Sisi, Seif, Alireza, & Jiang, Liang. "Quantum advantages for Pauli channel estimation." Phys. Rev. A 105.3 (2022): 032435.

Learning Hamiltonians from Long Time Evolutions

Consider the task of learning the coefficients of an n-qubit Hamiltonian H to ε error in ℓ_∞. Let $H = \sum_{a=1}^{m} \lambda_a E_a$ for $-1 \leq \lambda_a \leq 1$ and E_a a tensor product of Paulis with support size $O(1)$ and which is local with respect to some underlying lattice.

Suppose we want to learn some coefficient λ_a.

Problem

Can we do this, when only given the ability to perform the channel $\rho \mapsto e^{-iHt}\rho e^{iHt}$ for all $t > 100$? In other words, can we perform time-efficient Hamiltonian learning with time resolution beyond a small constant? Can we get it with total evolution time $\log(n)/\varepsilon$?

Total evolution time $1/\varepsilon$ and time resolution $\Theta(1)$ is known (Bakshi et al. '24), but nothing beyond this.

- Huang, Hsin-Yuan, Tong, Yu, Fang, Di, & Su, Yuan. "Learning many-body Hamiltonians with Heisenberg-limited scaling." Phys. Rev. Lett. 130.20 (2023): 200403.
- Dutkiewicz, Alicja, O'Brien, Thomas E., & Schuster, Thomas. "The advantage of quantum control in many-body Hamiltonian learning." arXiv:2304.07172 (2023).
- Bakshi, Ainesh, Liu, Allen, Moitra, Ankur, & Tang, Ewin. "Structure learning of Hamiltonians from real-time evolution." arXiv:2405.00082 (2024).

Lower Bounds for Hamiltonian Learning from Real-Time Evolution

Problem

Does Hamiltonian learning require a total time evolution dependent on n?

The best lower bound is $1/\varepsilon$, which follows from a simple "hybrid" argument: perturbing a coefficient by ε only affects the algorithm by $O(\varepsilon T)$ where T is the total evolution time, so T has to be large enough to detect this difference.
However, not even a $\log(n)/\varepsilon$ lower bound is known here: note that "union bound" style arguments do not work, since it is possible in simple examples to learn a Hamiltonian with probability 1. But the best algorithms only get $\Theta(\log(n)/\varepsilon)$, precisely from a union bound.
Further, we might expect the algorithm to require a dependence on one-spin energy $\max_{i\in[n]} \sum_{a\in[m]:\mathrm{supp}(E_a)\ni i} \|\lambda_a E_a\|$.

Bakshi, Ainesh, Liu, Allen, Moitra, Ankur, & Tang, Ewin. "Structure learning of Hamiltonians from real-time evolution." arXiv:2405.00082 (2024).

Does Amplitude Amplification Require Inverses?

Grover-style methods typically require the use of the inverse. For example, to estimate an entry of a unitary U, a naive approach requires $1/\varepsilon^2$ uses of U, but if one has access to $U^\dagger$, one can use amplitude estimation, which only requires $1/\varepsilon$ uses (of U and $U^\dagger$).

Problem

Can we prove that this use of $U^\dagger$ is somehow necessary?

Note that it is possible to convert d^3 uses of U to one use of $U^\dagger$ via Quintino et al., but this is extremely expensive as the number of qubits grows. So, can we show a separation between having the inverse and not for amplitude estimation, in the regime where d is sufficiently large? This question was previously posed by Apeldoorn et al. '23.

Quintino, Marco Túlio, Dong, Qingxiuxiong, Shimbo, Atsushi, Soeda, Akihito, & Murao, Mio. "Reversing unknown quantum transformations: universal quantum circuit for inverting general unitary operations." Phys. Rev. Lett. 123.21 (2019): 210502.

Apeldoorn, Joran van, Cornelissen, Arjan, Gilyén, András, & Nannicini, Giacomo. "Quantum tomography using state-preparation unitaries." SODA (2023): 1265-1318.

Polemical Discussion

In the last 5 years, there has been enormous progress in the following directions:

1. We now have efficient quantum algorithms, even NISQ algorithms, for learning properties of quantum states, quantum Hamiltonians, and low-depth quantum circuits.
2. We now have large quantum query complexity separations between different models of quantum computation (classical vs. quantum measurement protocols, NISQ vs. BQP, etc.) for certain kinds of property testing problems.

The tools (1) are extraordinary useful, but 'blunt' instruments. They also require a fine degree of local control over a quantum system which is not available in experimental platforms. The results (2) are very technically and conceptually interesting, but while the problems in question are mathematically nice they are for the most part not practically useful.

Problem

How do we go beyond (1) and (2)? For instance, how could we use a quantum computer to search for materials which furnish high-temperature superconductivity?

Aharonov, Dorit, Cotler, Jordan, & Qi, Xiao-Liang. "Quantum algorithmic measurement." Nature Communications 13.1 (2022):