Title: Learning quantum objects

Speakers: Amira Abbas

Collection: Foundations of Quantum Computational Advantage

Date: May 01, 2024 - 10:00 AM

URL: https://pirsa.org/24050005

Abstract: Whilst tomography has dominated the theory behind reconstructing/approximating quantum objects, such as states or channels, conducting full tomography is often not necessary in practice. If one is interested in learning properties of a quantum system, side-stepping the exponential lower bounds of tomography is then possible. In this talk, we will introduce various learning models for approximating quantum objects, survey the literature of quantum learning theory and explore instances where learning can be fully time- and sample efficient.

Learning quantum objects

Amira Abbas

University of Amsterdam, QuSoft

Foundations of Quantum Computational Advantage

May 2024

QuSoft
Research Center for Quantum Software

Universiteit van Amsterdam

Motivation: Approximate learnir

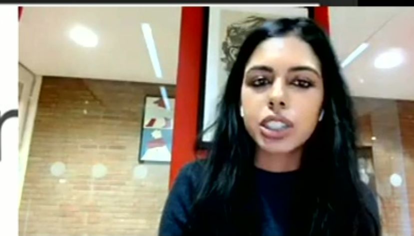

- Learning = approximating a function/object from samples

$$\mathbb{P}_{x\sim\mathcal{D}}[|f(x) - h(x)| \leq \epsilon] \geq 1 - \delta$$

Motivation: Approximate learnir

- Learning = approximating a function/object from samples
- Exact learning is often hard (need to see many examples)
- Very natural scenario (nature)

Motivation: Approximate learnir

- Learning = approximating a function/object from samples
- Exact learning is often hard (need to see many examples)
- Very natural scenario (nature)
- Interested in sample and time complexity of learning

$$\mathbb{P}_{x\sim\mathcal{D}}[|f(x) - h(x)| \leq \epsilon] \geq 1 - \delta$$

Motivation: Approximate learnir

- Learning = approximating a function/object from samples
- Exact learning is often hard (need to see many examples)
- Very natural scenario (nature)
- Interested in sample and time complexity of learning
  - Hardness of learning certain objects
  - Informs us on algorithmic design
- Two frameworks that dominate classical learning literature:
  - PAC Learning
  - Agnostic Learning

PAC/Agnostic Learning

Samples	$x \sim \mathcal{D}$
Target object	$f : \mathbb{F}_2^n \to \{-1, 1\}$
Concept class	$f \in \mathcal{C}$
Model/algo	$h \in \mathcal{H}$

PAC/Agnostic Learning

Goal: Algo must output a hypothesis function that is "close" to the target function

$$\mathbb{P}_{x\sim\mathcal{D}}[h(x) \neq f(x)] \leq \epsilon$$

Given access to samples $\{(x_i, f(x_i))\}_{i=1}^{m}$

$$m = \Theta(2^n/\epsilon)$$

(Boolean)

Hanneke, Steve. "The optimal sample complexity of PAC learning." Journal of Machine Learning Research 17, no. 38 (2016): 1-15.

Learning quantum states

Full state tomography

Given copies (or "samples") $\rho^{\otimes m}$

Output σ s.t.

$$\|\rho - \sigma\|_{\mathrm{Tr}} \leq \epsilon$$

$$m = \Theta(2^{2n})$$

O'Donnell, Ryan, and John Wright. "Efficient quantum tomography." In Proceedings of the forty-eighth annual ACM symposium on Theory of Computing, pp. 899-912. 2016.

Learning quantum states

Given samples: $\rho^{\otimes m}$ and promised that it is a stabiliser state, $\rho = |\phi\rangle\langle\phi|$

$$m = \text{poly}(n)$$

$$\text{Time} = \text{poly}(n)$$

Rocchetto, Andrea. "Stabiliser states are efficiently PAC-learnable." arXiv preprint arXiv:1705.00345 (2017).

Learning quantum states

Given samples: $\rho^{\otimes m}$ and promised that it is a stabiliser state, $\rho = |\phi\rangle\langle\phi|$

Lemma 1. *Let $E = (I + P)/2$ be a POVM measurement associated to a Pauli operator P and ρ an n-qubit stabiliser state then $\mathrm{Tr}(E\rho)$ can only take the following values $\{0, 1/2, 1\}$ and:*

$$\begin{cases} \textit{if } \mathrm{Tr}(E\rho) = 1 \textit{ then } P \textit{ is a stabiliser of } \rho; \\ \textit{if } \mathrm{Tr}(E\rho) = 1/2 \textit{ then neither } P \textit{ nor } -P \textit{ is a stabiliser of } \rho; \\ \textit{if } \mathrm{Tr}(E\rho) = 0 \textit{ then } -P \textit{ is a stabiliser of } \rho. \end{cases}$$

Rocchetto, Andrea. "Stabiliser states are efficiently PAC-learnable." arXiv preprint arXiv:1705.00345 (2017).

Learning quantum states

Given samples: $\rho^{\otimes m}$ and promised that it is a stabiliser state, $\rho = |\phi\rangle\langle\phi|$

Lemma 1. *Let $E = (I + P)/2$ be a POVM measurement associated to a Pauli operator P and ρ an n-qubit stabiliser state then $\mathrm{Tr}(E\rho)$ can only take the following values $\{0, 1/2, 1\}$ and:*

$$\begin{cases} \textit{if } \mathrm{Tr}(E\rho) = 1 \textit{ then } P \textit{ is a stabiliser of } \rho; \\ \textit{if } \mathrm{Tr}(E\rho) = 1/2 \textit{ then neither } P \textit{ nor } -P \textit{ is a stabiliser of } \rho; \\ \textit{if } \mathrm{Tr}(E\rho) = 0 \textit{ then } -P \textit{ is a stabiliser of } \rho. \end{cases}$$

With access to samples, $(E_i, \mathrm{Tr}[E_i\rho])$, the "labels" tell us if P_i is in stab$(|\phi\rangle)$

With enough samples, Gaussian elimination to find enough indep generators

Rocchetto, Andrea. "Stabiliser states are efficiently PAC-learnable." arXiv preprint arXiv:1705.00345 (2017).

Learning quantum states

Given samples: $\rho^{\otimes m}$ and a set of *known* states $S = \{\sigma_i\}$, output a $\sigma_j \in S$

s.t.

$$\|\rho - \sigma_j\|_{\mathrm{Tr}} \approx \min_{\sigma \in S} \|\rho - \sigma\|_{\mathrm{Tr}}$$

If S is the set of all possible stabiliser states, $S = \{\sigma_i\}_{i=1}^{2^{n^2}}$

$$m = O(n^2/\epsilon^2)$$

Huang, Hsin-Yuan, Richard Kueng, and John Preskill. "Predicting many properties of a quantum system from very few measurements." Nature Physics 16, no. 10 (2020): 1050-1057.

Learning quantum states

Given samples: $\rho^{\otimes m}$ and a set of *known* states $S = \{\sigma_i\}$, output a $\sigma_j \in S$

s.t.

$$\|\rho - \sigma_j\|_{\mathrm{Tr}} \approx \min_{\sigma \in S} \|\rho - \sigma\|_{\mathrm{Tr}}$$

If S is the set of all possible stabiliser states, $S = \{\sigma_i\}_{i=1}^{2^{n^2}}$

$$\text{Time} \sim \text{poly}(n) \qquad \text{(promise)}$$

Grewal, S., Iyer, V., Kretschmer, W., & Liang, D. (2023). Improved stabilizer estimation via bell difference sampling. arXiv preprint arXiv:2304.13915.

Learning quantum channels

S is the set of unital channels?

Given $\mathcal{N}^{\otimes m}$ and a set of known channels $S = \{\mathcal{U}_i\}$, output a $\mathcal{U}_j \in S$ s.t.

$$\|\mathcal{N} - \mathcal{U}_j\|_\diamond \leq \min_{\mathcal{U} \in S} \|\mathcal{N} - \mathcal{U}\|_\diamond + \epsilon$$

Cheng, Hao-Chung, Nilanjana Datta, Nana Liu, Theshani Nuradha, Robert Salzmann, and Mark M. Wilde. "Sample complexity of quantum hypothesis testing." arXiv preprint arXiv:2403.17868 (2024).

Learning quantum states

$$\|\rho - \sigma\|_{\mathrm{Tr}} \leq \epsilon \qquad m = \Theta(2^{2n})$$

$$\|\rho - \sigma_j\|_{\mathrm{Tr}} \approx \min_{\sigma \in S} \|\rho - \sigma\|_{\mathrm{Tr}} \qquad m = O(n^2/\epsilon^2)$$

What about S as the set of Gaussian states or Matchgate states?

Learning quantum channels

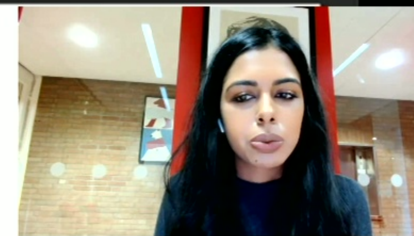

S is the set of unital channels?

Given $\mathcal{N}^{\otimes m}$ and a set of known channels $S = \{\mathcal{U}_i\}$, output a $\mathcal{U}_j \in S$ s.t.

$$\|\mathcal{N} - \mathcal{U}_j\|_\diamond \leq \min_{\mathcal{U} \in S} \|\mathcal{N} - \mathcal{U}\|_\diamond + \epsilon$$

Cheng, Hao-Chung, Nilanjana Datta, Nana Liu, Theshani Nuradha, Robert Salzmann, and Mark M. Wilde. "Sample complexity of quantum hypothesis testing." arXiv preprint arXiv:2403.17868 (2024).

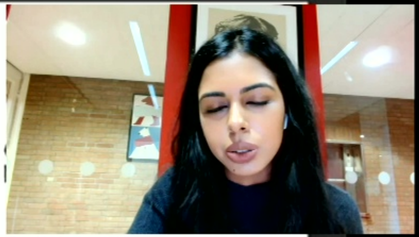

Learning *properties* of quantum

Full state tomography:

$$\max_E |\mathrm{Tr}[E\rho] - \mathrm{Tr}[E\sigma]| \leq \epsilon$$

Learning *properties* of quantum

Given samples: $\rho^{\otimes m}$ and $\{E_1, E_2, ..., E_M\}$, output σ s.t.

$$|\mathrm{Tr}[E_i \rho] - \mathrm{Tr}[E_i \sigma]| \leq \epsilon \quad \forall\ i$$

Shadow tomography

$$m = \mathrm{poly}(\log(M), n, 1/\epsilon)$$

$$\mathrm{Time} = \exp(n)$$

Aaronson, Scott. "Shadow tomography of quantum states." In Proceedings of the 50th annual ACM SIGACT symposium on theory of computing, pp. 325-338. 2018.

Learning *properties* of quantum

Given samples: $\rho^{\otimes m}$ and $\{P_1, P_2, ..., P_M\}$, output σ s.t.

$$|\mathrm{Tr}[P_i \rho] - \mathrm{Tr}[P_i \sigma]| \leq \epsilon \quad \forall\, i$$

Shadow tomography for local Pauli operators

$$m = \mathrm{poly}(\log(M), 1/\epsilon)$$

$$\mathrm{Time} = \mathrm{poly}(n)$$

Huang, Hsin-Yuan, Richard Kueng, and John Preskill. "Information-theoretic bounds on quantum advantage in machine learning." Physical Review Letters 126, no. 19 (2021): 190505.

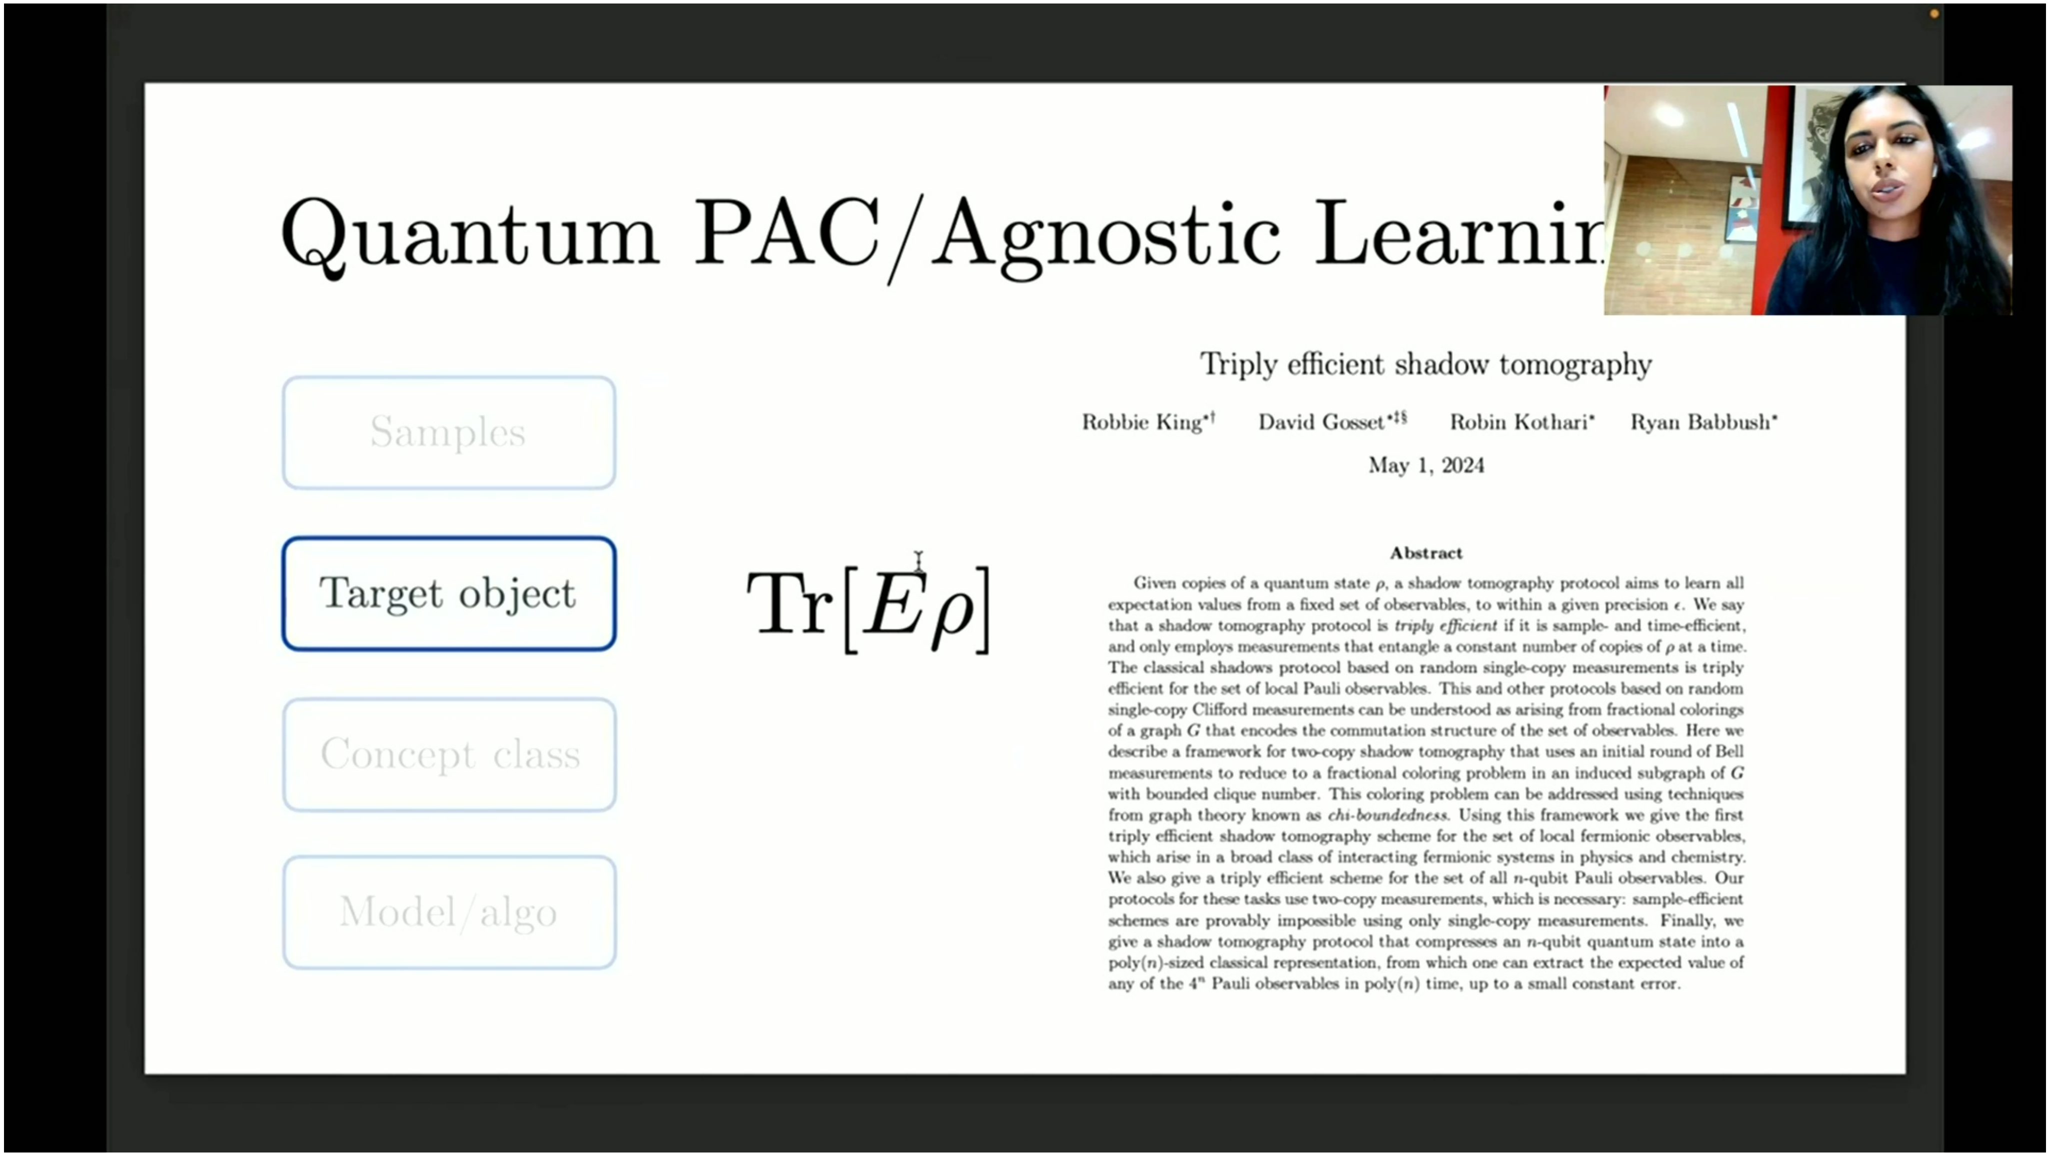
Quantum PAC/Agnostic Learnin
Samples
Target object
Concept class
Model/algo
Tr[Eρ]
Triply efficient shadow tomography
Robbie King*† David Gosset*‡§ Robin Kothari* Ryan Babbush*
May 1, 2024
Abstract
Given copies of a quantum state ρ, a shadow tomography protocol aims to learn all expectation values from a fixed set of observables, to within a given precision ϵ. We say that a shadow tomography protocol is triply efficient if it is sample- and time-efficient, and only employs measurements that entangle a constant number of copies of ρ at a time. The classical shadows protocol based on random single-copy measurements is triply efficient for the set of local Pauli observables. This and other protocols based on random single-copy Clifford measurements can be understood as arising from fractional colorings of a graph G that encodes the commutation structure of the set of observables. Here we describe a framework for two-copy shadow tomography that uses an initial round of Bell measurements to reduce to a fractional coloring problem in an induced subgraph of G with bounded clique number. This coloring problem can be addressed using techniques from graph theory known as chi-boundedness. Using this framework we give the first triply efficient shadow tomography scheme for the set of local fermionic observables, which arise in a broad class of interacting fermionic systems in physics and chemistry. We also give a triply efficient scheme for the set of all n-qubit Pauli observables. Our protocols for these tasks use two-copy measurements, which is necessary: sample-efficient schemes are provably impossible using only single-copy measurements. Finally, we give a shadow tomography protocol that compresses an n-qubit quantum state into a poly(n)-sized classical representation, from which one can extract the expected value of any of the 4^n Pauli observables in poly(n) time, up to a small constant error.

Conclusion:

- There is a lot of freedom in quantum learning theory
- Changing a learning task could enable efficiency, or hardness
- Duality of learning (cryptography)
- Many open questions in shadow tomography for states, quantum PAC learning, channel learning etc.
- There is room to exploit states with specific structure, which seems relatively unexplored

Thank you!

amira@cwi.nl