Learning Quantum Processes with Quantum Statistical Queries

Chirag Wadhwa, Mina Doosti

arxiv:2310.02075

THE UNIVERSITY *of* EDINBURGH

Outline

Quantum Statistical Query Learning

Average-case lower bounds for QPSQ learners

Learning Algorithm and Cryptographic Attack

Outline

Quantum Statistical Query Learning

Average-case lower bounds for QPSQ learners

Learning Algorithm and Cryptographic Attack

Background : Quantum Statistical Queries

- Classically, **Statistical Query** oracles were introduced to formalize learning from *aggregate properties* of data rather than individual samples [Kea98].

Background : Quantum Statistical Queries

- Classically, **Statistical Query** oracles were introduced to formalize learning from *aggregate properties* of data rather than individual samples [Kea98].
- Similarly, **Quantum Statistical Query** oracles formalize learning a state from *expectation values* [AGY20].

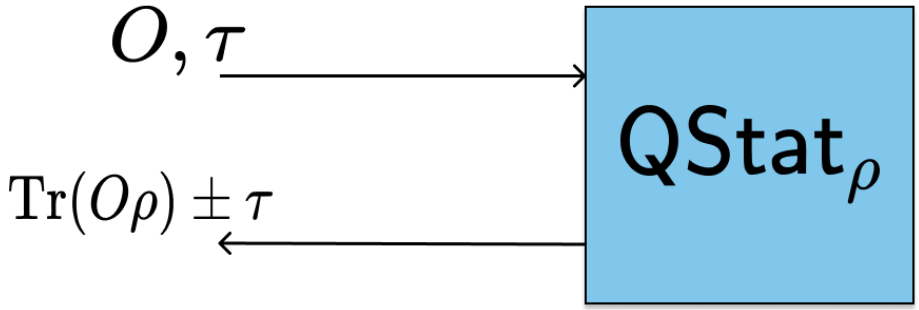

Quantum Statistical Queries to Processes

We define the QPStat oracle for learning a quantum process.[1]

[1] Defined independently by [NPVY24].

Quantum Statistical Queries to Processes

We define the QPStat oracle for learning a quantum process.[1]

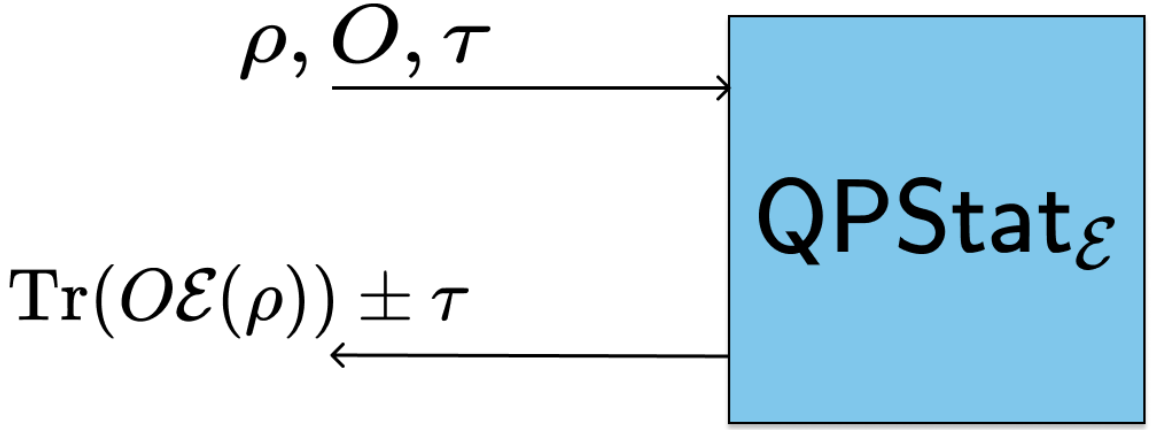

[1] Defined independently by [NPVY24].

Efficient QPSQ Learners

A learning algorithm is efficient if it has efficient query & time complexities. Efficient QPSQ algorithms also have the following properties.

Efficient QPSQ Learners

A learning algorithm is efficient if it has efficient query & time complexities. Efficient QPSQ algorithms also have the following properties.

- **Efficient Query Tolerance:** We require τ to not be too small. In general, $1/\tau \leq \text{poly(n)}$.

Efficient QPSQ Learners

A learning algorithm is efficient if it has efficient query & time complexities. Efficient QPSQ algorithms also have the following properties.

- **Efficient Query Tolerance:** We require τ to not be too small. In general, $1/\tau \leq \text{poly(n)}$.
- **Efficient preparation and measurement:** We require the queried ρ, O to be efficiently preparable and measurable.

Efficient QPSQ Learners

A learning algorithm is efficient if it has efficient query & time complexities. Efficient QPSQ algorithms also have the following properties.

- **Efficient Query Tolerance:** We require τ to not be too small. In general, $1/\tau \leq \text{poly(n)}$.
- **Efficient preparation and measurement:** We require the queried ρ, O to be efficiently preparable and measurable.

These requirements ensure that the oracle queries can be implemented efficiently in practice.

Learning problem

We look at learning a class of quantum channels from QPStat queries.

Learning problem

We look at learning a class of quantum channels from QPStat queries.

- Consider the class $\mathcal{C} = \{\mathcal{E}_i\}_i$.

Learning problem

We look at learning a class of quantum channels from QPStat queries.

- Consider the class $\mathcal{C} = \{\mathcal{E}_i\}_i$.
- Given access to $\mathrm{QPStat}_{\mathcal{E}}$, for an unknown $\mathcal{E} \in \mathcal{C}$, produce a hypothesis $\mathcal{E}'$, such that with high probability,

$$\mathrm{d}(\mathcal{E}, \mathcal{E}') \leq \epsilon.$$

Learning problem

We look at learning a class of quantum channels from QPStat queries.

- Consider the class $\mathcal{C} = \{\mathcal{E}_i\}_i$.
- Given access to $\mathrm{QPStat}_{\mathcal{E}}$, for an unknown $\mathcal{E} \in \mathcal{C}$, produce a hypothesis $\mathcal{E}'$, such that with high probability,

$$\mathrm{d}(\mathcal{E}, \mathcal{E}') \leq \epsilon.$$

- We will consider the diamond distance.

$$d(\mathcal{E}_1, \mathcal{E}_2) = \frac{\|\mathcal{E}_1 - \mathcal{E}_2\|_\diamond}{2}, \quad \|\mathcal{E}\|_\diamond = \max_\rho \|(\mathcal{E} \otimes \mathcal{I})\rho\|_1.$$

Outline

Quantum Statistical Query Learning

Average-case lower bounds for QPSQ learners

Learning Algorithm and Cryptographic Attack

Average-case lower bounds for 2-design/Haar-random unitaries

Average-case lower bounds for 2-design/Haar-random unitaries

Theorem (Lower bound for unitary 2-designs)

Any algorithm for learning a unitary 2-design with probability at least β over the random unitary must make

$$\Omega(\beta\tau^2 2^n)$$

QPStat *queries of tolerance at least τ.*

Average-case lower bounds for 2-design/Haar-random unitaries

Theorem (Lower bound for unitary 2-designs)

Any algorithm for learning a unitary 2-design with probability at least β over the random unitary must make

$$\Omega(\beta\tau^2 2^n)$$

QPStat *queries of tolerance at least τ.*

Theorem (Lower bound for Haar-random unitaries)

Any algorithm for learning a Haar-random unitary with probability at least β over the random unitary must make

$$\beta \exp\left(\tilde{\Omega}(2^n\tau^2)\right)$$

QPStat *queries of tolerance at least τ.*

Hardness of distinguishing

- Reduction to a many-vs-one distinguishing problem[2].

[2] [Fel17], [NIS+23],[AHS24],[Nie24].

Hardness of distinguishing

- Reduction to a many-vs-one distinguishing problem[2].

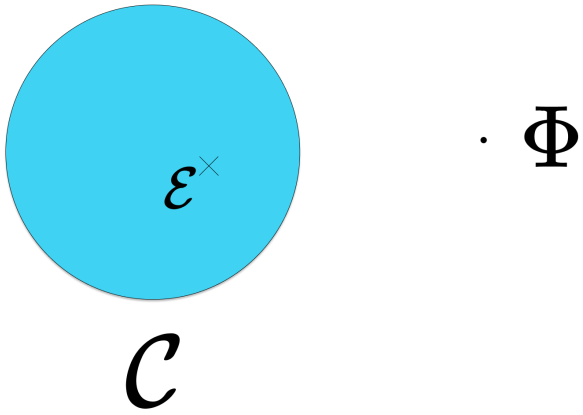

[2] [Fel17], [NIS+23],[AHS24],[Nie24].

Hardness of distinguishing

- Reduction to a many-vs-one distinguishing problem[2].

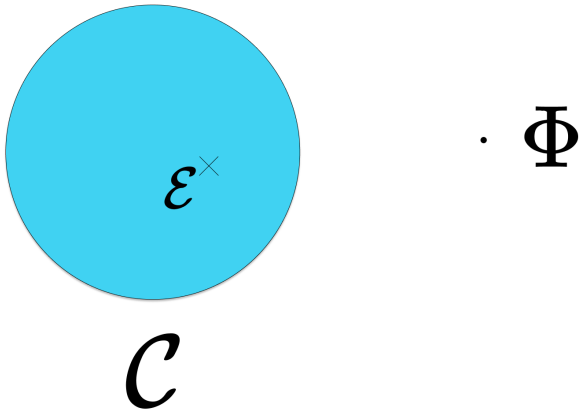

- The probability that a single QPStat query (ρ, O, τ) cannot distinguish between $\mathcal{C}$ and Φ:

$$\mathbf{Pr}_{\mathcal{E}\sim\mathcal{C}}[\mathrm{Tr}(O\mathcal{E}(\rho)) \approx_\tau \mathrm{Tr}(O\Phi(\rho))]$$

[2] [Fel17], [NIS+23],[AHS24],[Nie24].

Hardness of distinguishing

- Reduction to a many-vs-one distinguishing problem[2].

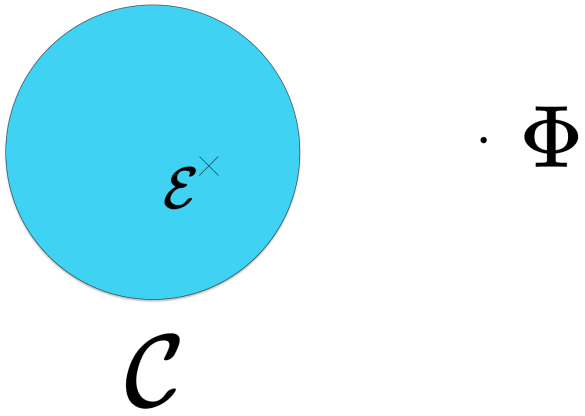

- The probability that a single QPStat query (ρ, O, τ) cannot distinguish between $\mathcal{C}$ and Φ:

$$\mathbf{Pr}_{\mathcal{E}\sim\mathcal{C}}[\mathrm{Tr}(O\mathcal{E}(\rho)) \approx_\tau \mathrm{Tr}(O\Phi(\rho))]$$

- If this probability is large, many queries are needed to solve the distinguishing task!

[2] [Fel17], [NIS+23],[AHS24],[Nie24].

Triviality

Even when the distinguishing task is hard, the learning problem can be trivial.

Triviality

Even when the distinguishing task is hard, the learning problem can be trivial.

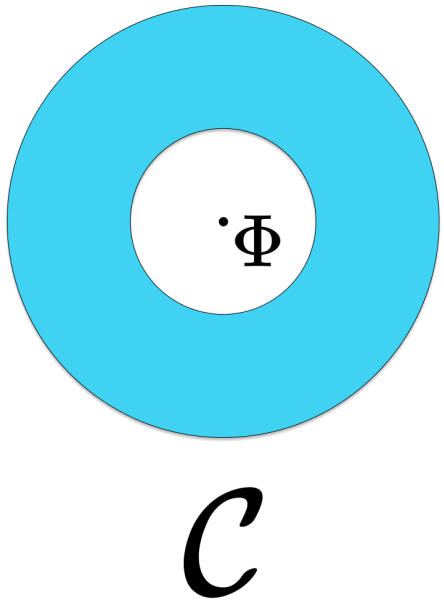

Triviality

Even when the distinguishing task is hard, the learning problem can be trivial.

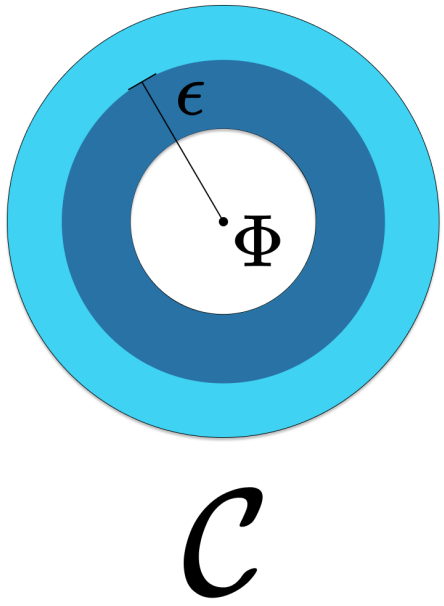

Lower Bounds

For meaningful lower bounds, choose Φ:

Lower Bounds

For meaningful lower bounds, choose Φ:

- The outputs of QPStat queries are indistinguishable.

Lower Bounds

For meaningful lower bounds, choose Φ:

- The outputs of QPStat queries are indistinguishable.
- Φ is far from the target class.

Lower Bounds

For meaningful lower bounds, choose Φ:

- The outputs of QPStat queries are indistinguishable.
- Φ is far from the target class.

For unitary 2-designs and Haar-random unitaries, this is satisfied by the maximally depolarizing channel $\Phi_{dep} : \rho \to \mathrm{Tr}(\rho)\frac{\mathbb{I}}{2^n}$.

Lower Bounds

For meaningful lower bounds, choose Φ:

- The outputs of QPStat queries are indistinguishable.
- Φ is far from the target class.

For unitary 2-designs and Haar-random unitaries, this is satisfied by the maximally depolarizing channel $\Phi_{dep} : \rho \to \mathrm{Tr}(\rho)\frac{\mathbb{I}}{2^n}$.

- The output expectation values of 2-design/Haar-random unitaries concentrate around those of Φ_{dep}.

Lower Bounds

For meaningful lower bounds, choose Φ:

- The outputs of QPStat queries are indistinguishable.
- Φ is far from the target class.

For unitary 2-designs and Haar-random unitaries, this is satisfied by the maximally depolarizing channel $\Phi_{dep} : \rho \to \mathrm{Tr}(\rho)\frac{\mathbb{I}}{2^n}$.

- The output expectation values of 2-design/Haar-random unitaries concentrate around those of Φ_{dep}.
- All unitaries are far from Φ_{dep}.

Outline

Quantum Statistical Query Learning

Average-case lower bounds for QPSQ learners

Learning Algorithm and Cryptographic Attack

Predicting expectation values of an unknown quantum process

- For a known observable O, given QPStat access to an unknown channel $\mathcal{E}$, we want to predict $\mathrm{Tr}(O\mathcal{E}(\rho))$ for new states ρ.

Predicting expectation values of an unknown quantum process

- For a known observable O, given QPStat access to an unknown channel $\mathcal{E}$, we want to predict $\mathrm{Tr}(O\mathcal{E}(\rho))$ for new states ρ.

Theorem

For a known local observable O and certain distributions D over quantum states, given QPStat access to an unknown quantum channel $\mathcal{E}$, there exists an efficient learning algorithm to predict $\mathrm{Tr}(O\mathcal{E}(\rho))$ with low error on average over $\rho \sim D$.

Predicting expectation values of an unknown quantum process

- For a known observable O, given QPStat access to an unknown channel $\mathcal{E}$, we want to predict $\mathrm{Tr}(O\mathcal{E}(\rho))$ for new states ρ.

Theorem

For a known local observable O and certain distributions D over quantum states, given QPStat access to an unknown quantum channel $\mathcal{E}$, there exists an efficient learning algorithm to predict $\mathrm{Tr}(O\mathcal{E}(\rho))$ with low error on average over $\rho \sim D$.

- Extend the algorithm of [HCP23] to the QPStat setting.

Classical Readout QPUFs

Classical Readout QPUFs

- Quantum Physical Unclonable Functions (QPUFs): A hardware primitive providing secure authentication **without computational assumptions**.

Classical Readout QPUFs

- Quantum Physical Unclonable Functions (QPUFs): A hardware primitive providing secure authentication **without computational assumptions**.
- Classical Readout QPUFs: A lightweight version of QPUFs, but not known to be secure.

Classical Readout QPUFs

- Quantum Physical Unclonable Functions (QPUFs): A hardware primitive providing secure authentication **without computational assumptions**.
- Classical Readout QPUFs: A lightweight version of QPUFs, but not known to be secure.
- The behaviour of a CR-QPUF can be modelled by a QPStat oracle.

Classical Readout QPUFs

- Quantum Physical Unclonable Functions (QPUFs): A hardware primitive providing secure authentication **without computational assumptions**.
- Classical Readout QPUFs: A lightweight version of QPUFs, but not known to be secure.
- The behaviour of a CR-QPUF can be modelled by a QPStat oracle.
- Our learning algorithm for predicting expectation values from QPStat queries can be used to attack CR-QPUFs!

Summary

Summary

- Defined quantum statistical queries for learning quantum processes.

Summary

- Defined quantum statistical queries for learning quantum processes.
- Lower bounds for unitary 2-designs and Haar-random unitaries.

Summary

- Defined quantum statistical queries for learning quantum processes.
- Lower bounds for unitary 2-designs and Haar-random unitaries.
- Learning algorithm to predict expectation values of arbitrary processes.

Summary

- Defined quantum statistical queries for learning quantum processes.
- Lower bounds for unitary 2-designs and Haar-random unitaries.
- Learning algorithm to predict expectation values of arbitrary processes.
- Attack on Classical Readout QPUFs in most practical regimes.

Open Questions

Open Questions

- Tradeoffs between restrictions on observables, states and channels for predicting expectation values?

Open Questions

- Tradeoffs between restrictions on observables, states and channels for predicting expectation values?
- Meaningful separations between classical and QPSQ learners for learning classical functions?