Online Learning of a Panoply of Quantum Objects

Akshay Bansal[1], Ian George[2], Soumik Ghosh[3], Jamie Sikora[1], Alice Zheng[1]

[1]Virginia Tech, [2]UIUC, [3]UChicago

Introduction

When utilizing quantum algorithms, accurately learning their outputs is crucial. Given physical quantum objects such algorithms produce, we wish to learn their mathematical description. Since measurements are destructive, doing so consumes the objects – an expended resource that we wish to minimize. This setting has been extensively studied for quantum states and channels, and we aim to *generalize* one possible approach to other objects. Our extended setting allows the sets of objects and co-objects to be arbitrary convex subsets of positive semidefinite operators, thus encompassing a broad range of quantum objects.

In this setting, we learn a quantum object via its interactions with numerous co-objects (e.g. learn a quantum state by its interactions with measurements). At each time step, the adversary queries one such co-object and the learner responds with an object, being penalized through a loss function evaluated at the inner product of the two. We aim to minimize *regret*

$$\mathcal{R}_T = \sum_{t=1}^{T} f_t(\omega_t) - \min_{\varphi \in \mathcal{K}} \sum_{t=1}^{T} f_t(\varphi) \ ,$$

that is, the accumulated expected loss of an algorithm compared to the best fixed action.

Methods

We extend a procedure for online shadow tomography based on RFTL (Regularized Follow-the-Leader). We use methods from online convex optimization and complex matrix analysis to extend the relevant proofs.

Input: T, $\eta > 0$, convex and compact $\mathcal{K} \subseteq \mathrm{Pos}(\mathcal{X})$, bounded $\mathcal{E} \subseteq \mathrm{Herm}(\mathcal{X})$.

Set initial hypothesis $\omega_1 \leftarrow \arg\max_{\varphi\in\mathcal{K}}\{S(\varphi)\}$, where $S(X) = -\mathrm{Tr}(X \ln(X))$.

for $t \leftarrow 1$ **to** T **do**

- Predict ω_t and incur cost $f_t(\omega_t) := \ell_t(\langle E_t, \omega_t\rangle)$ with $\ell_t : \mathcal{R} \to \mathcal{R}$ and $E_t \in \mathcal{E}$.
- Let $\ell_t'(x)$ be a sub-derivative of ℓ_t with respect to x and define
$$\nabla_t \leftarrow \ell_t'(\langle E_t, \omega_t\rangle) E_t \ .$$
- Update decision according to the RFTL rule
$$\omega_{t+1} \leftarrow \arg\min_{\varphi\in\mathcal{K}} \left\{ \eta \sum_{s=1}^{t} \langle \nabla_s, \varphi\rangle - S(\varphi) \right\} \ .$$

end

We establish the *learnability* of many classes of quantum objects. Given a set of objects $\mathcal{K}$ with trace bounded by A and a set of co-objects $\mathcal{E}$ with operator norm bounded by C, the algorithm guarantees

$$\mathcal{R}_T \leq 4BCD\sqrt{AT} \ .$$

Here, the loss functions are B-Lipschitz and D is the diameter of $\mathcal{K}$ relative to quantum entropy. Additionally, we have established a *generalized version of Pinsker's inequality* applicable to unnormalized distributions

$$\frac{1}{4}\|P-Q\|_{\mathrm{Tr}}^2 \leq \max\{\mathrm{Tr}(P), \mathrm{Tr}(Q)\}\,[D(P\,\|\,Q) - \mathrm{Tr}(P-Q)] \ .$$

The above constant may be tightened to $1/2$ via the use of a strong convexity-based argument.

Applications

Examples of learnable objects:

States $\mathcal{R}_T \leq \left(4B\sqrt{\ln(\dim(\mathcal{X}))}\right)\sqrt{T}$

Gram matrices $\mathcal{R}_T \leq \left(4Bn\sqrt{\ln(n)}\right)\sqrt{T}$

Effects $\mathcal{R}_T \leq \left(4B\dim(\mathcal{X})\sqrt{e^{-1} + \ln(\dim(\mathcal{X}))}\right)\sqrt{T}$

Channels $\mathcal{R}_T \leq \left(4B\ \dim(\mathcal{X})\sqrt{\ln(\dim(\mathcal{X})) + \ln(\dim(\mathcal{Y}))}\right)\sqrt{T}$

Interactive measurements

$$\mathcal{R}_T \leq \left(4B\ \dim(\mathcal{X})\dim(\mathcal{Y})\sqrt{\ln(\dim(\mathcal{X})) + \ln(\dim(\mathcal{Y}))}\right)\sqrt{T}$$

Strategies

$$\mathcal{R}_T \leq \left(4B\left(\prod_{i=1}^{n}\dim(\mathcal{Y}_i)\right)\left(\prod_{i=1}^{n}\dim(\mathcal{X}_i)\right)\sqrt{\sum_{i=1}^{n}\ln(\dim(\mathcal{X}_i)) + \sum_{i=1}^{n}\ln(\dim(\mathcal{Y}_i))}\right)\sqrt{T}$$

Contributions

1. A sublinear regret bound for online learning of many classes of quantum objects.
2. A generalization of Pinsker's inequality for arbitrary PSD operators with possibly different traces, which could be of independent interest and applicable to more general classes of divergences.

Acknowledgements

We thank Asad Raza, Matthias C. Caro, Jens Eisert, and Sumeet Khatri for sharing a draft of their independent and concurrent work Online learning of quantum processes. I.G. was supported by the National Science Foundation under Grant No. 2112890 and in part by an Illinois Distinguished Fellowship. A.B. was partially supported by the BitShares Fellowship at Virginia Tech. This research was funded in part by the Commonwealth of Virginia's Commonwealth Cyber Initiative under Grant No. 467714.

Full paper and references:

2406.04245