Online learning of a panoply of quantum objects

arXiv:2406.04245

Akshay Bansal, Ian George, Soumik Ghosh, Jamie Sikora, Alice Zheng

Virginia Tech, UChicago, UIUC

Table of contents

Tomography

Quantum objects

- Subsets of positive semidefinite operators.

$$D(\mathcal{X}) := \{\rho \in \mathrm{Pos}(\mathcal{X}) : \mathrm{Tr}(\rho) = 1\} \ ,$$
$$C(\mathcal{X},\mathcal{Y}) := \{J \in \mathrm{Pos}(\mathcal{Y} \otimes \mathcal{X}) : \mathrm{Tr}_{\mathcal{Y}}(J) = \mathbb{1}_{\mathcal{X}}\} \ .$$

- Exponential number of parameters
  - n qubits $\rightarrow$ dimension $D = 2^n$ operator.
- Physical objects

Tomography

- How many physical copies of an object to learn its description?
    1. Precision (w.r.t. some norm)
    2. Destructive nature of measurements
- Quantum states: $O(D^2)$, but also $\Omega(D^2)$.
- Are the amplitudes *really* there?

Shadow Tomography

Changing the task

- Shadow tomography [Aar18].
  - *Input:* many same quantum *objects* ρ
  - *Input:* M quantum co-objects E_i
  - *Output:* b_i such that $b_i \approx \langle E_i, \rho \rangle$
- Knowing $b_i \to$ reasonably good info about ρ.
- Quantum money, advice, copy-protected software, etc.
- Simple approaches (quantum states):
  - Full tomography: $O(D^2)$
  - Just do the measurements: $O(M)$

One approach [Aar18]

- Online learning: iteratively refine hypothesis, given challenging measurements
  - *Gentle search*: find challenging measurements
- Quantum states:

$$\tilde{O}(\ln(D) \cdot \epsilon^{-2}) \times \tilde{O}(\ln^4(M) \cdot \epsilon^{-2}) = \tilde{O}(\ln(D) \cdot \ln^4(M) \cdot \epsilon^{-4}) \ .$$

- Upper bound on the number of *mistakes*.

Online Learning

Online Convex Optimization

- First defined in **machine learning** literature.
- Metric of success borrowed from **game theory**.
- Framework closely tied to **statistical learning theory** and **convex optimization**.

Online Convex Optimization

Convex optimization: single interaction

1. Receive convex set $\mathcal{K}$ and convex function f.
2. Choose $\mathbf{x} \in \mathcal{K}$ and suffer loss $f(\mathbf{x})$.

Online convex optimization: iterative process

1. Choose $\mathbf{x}_t \in \mathcal{K}$ and suffer loss $f_t(\mathbf{x}_t)$.
2. Receive convex loss function f_t.
3. Repeat.

- Choose $\mathbf{x}_t$ before f_t is known.
- Losses f_t can be different in general.
- Losses f_t can be selected **adversarially**.

Online Convex Optimization - Assumptions

Loss functions f_t are

- Bounded
- Convex

Decision set $\mathcal{K}$ is

- Bounded/structured
- Convex

Unbounded f_t

$$f_1(x) = \begin{cases} 0 & x \neq \mathbf{x}_t \ , \\ \infty & \text{otherwise} \ . \end{cases}$$

Unbounded $\mathcal{K}$

Assign loss 1 to all $\mathbf{x}_t$, while setting aside some strategies with 0 loss. Since $\mathcal{K}$ is infinite, the latter will never run out.

Regret

- Adversary can still force the learner to incur constant loss at every round.
- **Metric:** total accumulated loss *versus best fixed action*.

$$\mathcal{R}_T^{\mathcal{A}} = \sup_{f_1,\ldots,f_T \subseteq \mathcal{F}} \left\{ \sum_{t=1}^{T} f_t(\mathbf{x}_t^{\mathcal{A}}) - \min_{\mathbf{x} \in \mathcal{K}} \sum_{t=1}^{T} f_t(\mathbf{x}) \right\} .$$

- $\mathcal{A}$ and $\sup$ are typically clear from context and omitted.
- Interested in an **upper bound** for the **worst-case** regret.
- **Goal:** sublinear regret bound.

Applications

- Prediction from expert advice (experts problem)
- Online spam filtering
- Online shortest paths
- Recommendation systems
- Portfolio selection
- etc...

General algorithm

Algorithm 1 Regularized Follow-the-Leader (RFTL/FTRL)

Input: T, $\eta > 0$, convex regularization function R, a convex and compact set $\mathcal{K} \subseteq \mathrm{Herm}\,\mathcal{X}$.

1: Set initial hypothesis $\omega_1 \leftarrow \arg\min_{\varphi \in \mathcal{K}} \{R(\varphi)\}$.

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

3: Predict ω_t and incur cost $f_t(\omega_t)$, where $f_t : \mathcal{K} \to \mathbb{R}$ is convex.

4: Let ∇_t be a subgradient of f_t at ω_t.

5: Update decision according to the RFTL rule

$$\omega_{t+1} \leftarrow \arg\min_{\varphi \in \mathcal{K}} \left\{ \eta \sum_{s=1}^{t} \langle \nabla_s, \varphi \rangle + R(\varphi) \right\} .$$

6: **end for**

Regret bound: quantum states

Theorem ([ACH+18] Theorems 2 and 3)

Suppose that

- *Loss functions are B-Lipschitz*

RFTL can be applied to n-qubit quantum state tomography for a regret bound of

$$\mathcal{R}_T \leq 2B\sqrt{(2\ln 2)nT}$$

Results

Applied algorithm

Algorithm 2 RFTL for Online Learning of Quantum Objects

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

Set initial hypothesis $\omega_1 \leftarrow \arg\max_{\varphi \in \mathcal{K}} \{S(\varphi)\}$.

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 $\ell_t : \mathbb{R} \to \mathbb{R}$ 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 for

Regret bound: quantum objects

Theorem

Suppose that

- *Trace of objects is bounded by A*
- *Loss functions are B-Lipschitz*
- *Operator norm of co-objects is bounded by C*
- *The diameter of $\mathcal{K}$ w.r.t. S is* $D^2 := \max_{\varphi,\varphi' \in \mathcal{K}} \{S(\varphi') - S(\varphi)\}$.
- *There exists $\alpha > 0$ such that $\alpha \mathbb{1} \in \mathcal{K}$*

RFTL for Online Learning of Quantum States achieves a regret bound of

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

Bounding the diameter

Lemma

Suppose $\mathcal{K} \in \mathrm{Pos}(\mathcal{X})$ *satisfies* $\mathrm{Tr}(X) = A$ *for all* $X \in \mathcal{K}$ *and* $A \geq 1$. *Then,*

$$D^2 \leq A \ln(\dim(\mathcal{X})) \, .$$

If we have the same with $\mathrm{Tr}(X) \leq A$, *then*

$$D^2 \leq \begin{cases} A \ln(\dim(\mathcal{X})) & \text{if } A \leq e^{-1} \dim(\mathcal{X}) \, , \\ e^{-1} \dim(\mathcal{X}) + A \ln(A) & \text{if } A \geq e^{-1} \dim(\mathcal{X}) \, . \end{cases}$$

Example corollaries

Objects	Co-objects	Regret bound
States	Effects	$4B\sqrt{\ln(\dim(\mathcal{X}))}\sqrt{T}$
Effects	States	$4B\dim(\mathcal{X})\sqrt{e^{-1}+\ln(\dim(\mathcal{X}))}\sqrt{T}$
Gram matrices	Unit ball	$4Bn\sqrt{\ln(n)}\sqrt{T}$
Channels	Interactive msmts.	$4B\dim(\mathcal{X})\sqrt{\ln(\dim(\mathcal{X}))+\ln(\dim(\mathcal{Y}))}\sqrt{T}$
Interactive msmts.	Channels	$4B\dim(\mathcal{X}\mathcal{Y})\sqrt{\ln(\dim(\mathcal{X}))+\ln(\dim(\mathcal{Y}))}\sqrt{T}$
Strategies	Co-strategies	—"—
Co-strategies	Strategies	—"—

- For quantum strategies and co-strategies,

$$\dim(\mathcal{X}) = \prod_{i=1}^{n} \dim(\mathcal{X}_i)\ , \qquad \dim(\mathcal{Y}) = \prod_{i=1}^{n} \dim(\mathcal{Y}_i)\ .$$

Proof Sketch

Challenges

1. Variable traces
2. Derivatives of complex operators

Challenges

1. Variable traces
    - Generalized Pinsker's inequality
2. Derivatives of complex operators
    - Fréchet differentiation

(Quantum) Pinsker's inequality

Theorem

For any positive semidefinite P and Q with $\mathrm{Tr}(P) = \mathrm{Tr}(Q) = 1$,

$$\frac{1}{2}\|P - Q\|_{\mathrm{Tr}}^2 \leq D(P\|Q) \; .$$

(Quantum) Pinsker's inequality

Theorem

For any positive semidefinite P and Q with $\mathrm{Tr}(P) = \mathrm{Tr}(Q) > 0$,

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

- Relax $\mathrm{Tr}(P) = 1$ constraint via scaling

(Quantum) Pinsker's inequality

Theorem

For any positive semidefinite P and Q,

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

- Not as tight as we need...

Generalized Pinsker's inequality

Theorem

For any positive semidefinite P and Q,

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

Generalized Pinsker's inequality

Theorem

For any positive semidefinite P and Q,

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

- Also a tightening of Klein's inequality $0 \leq D(P\|Q) - \mathrm{Tr}(P-Q)$

Alternatives

1. Strong convexity of unnormalized Von Neumann entropy [Yu13]
    - Tightens 1/4 constant to 1/2 in Generalized Pinsker's
    - Original proof applies more generally to other f-divergences

Alternatives

1. Strong convexity of unnormalized Von Neumann entropy [Yu13]
   - Tightens 1/4 constant to 1/2 in Generalized Pinsker's
   - Original proof applies more generally to other f-divergences
2. "Slack" dimension

$$\overline{\mathcal{K}} = \left\{ \begin{bmatrix} X & 0 \\ 0 & \alpha - \mathrm{Tr}(X) \end{bmatrix} : X \in \mathcal{K} \right\}, \qquad \overline{\mathcal{E}} = \left\{ \begin{bmatrix} E & 0 \\ 0 & 0 \end{bmatrix} : E \in \mathcal{E} \right\}$$

   - Permits the use of ordinary Pinsker's inequality
   - Diameters of $\mathcal{K}$ and $\overline{\mathcal{K}}$ in general differ

Regret bound

Lemma ([Haz16] Lemma 5.3)

RFTL guarantees

$$\mathcal{R}_T \leq \sum_{t=1}^{T} \langle \nabla_t, \omega_t - \omega_{t+1} \rangle + \frac{1}{\eta} D^2 \ .$$

Lemma

The function

$$\Phi_E(X) := \langle E, X \rangle + \mathrm{Tr}(X \ln(X))$$

is Fréchet-differentiable with

$$\nabla \Phi_E(X) = E + \mathbb{1} + \ln(X) \ .$$

Bounding the norm

We have $B_{\Phi_E} = B_R = B_{-D}$, and Generalized Pinsker's inequality on ω_t and ω_{t+1} gives

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathrm{Tr}}^2 \leq \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\}\, B_{-D}(\omega_t || \omega_{t+1}) \; .$$

Bounding the norm

We have $B_{\Phi_E} = B_R = B_{-D}$, and Generalized Pinsker's inequality on ω_t and ω_{t+1} gives

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathrm{Tr}}^2 \leq \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\}\, B_{-D}(\omega_t||\omega_{t+1}) \ .$$

RFTL guarantees

$$\begin{aligned}
B_{-D}(\omega_t||\omega_{t+1}) &= \Phi_t(\omega_t) - \Phi_t(\omega_{t+1}) - \langle \nabla\Phi_t(\omega_{t+1}, \omega_t - \omega_{t+1}\rangle && \text{(Definition of } B_{\Phi_E}\text{)} \\
&\leq \Phi_t(\omega_t) - \Phi_t(\omega_{t+1}) && \text{(Variational inequality)} \\
&= \Phi_{t-1}(\omega_t) - \Phi_{t-1}(\omega_{t+1}) + \eta\langle \nabla_t, \omega_t - \omega_{t+1}\rangle && \text{(Definition of } \Phi_t\text{)} \\
&\leq \eta\langle \nabla_t, \omega_t - \omega_{t+1}\rangle \ . && \text{(Update rule)}
\end{aligned}$$

Bounding the inner product

Lemma

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}}^2 \leq \eta\langle\nabla_t, \omega_t - \omega_{t+1}\rangle$$

Bounding the inner product

Lemma

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}}^2 \leq \eta\langle\nabla_t, \omega_t - \omega_{t+1}\rangle$$

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}}^2 \leq \eta \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\} \langle\nabla_t, \omega_t - \omega_{t+1}\rangle \quad \text{(Combine with Gen. Pinsker's)}$$

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}} \leq 4\eta \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\}\|\nabla_t\|_{\mathrm{op}} \quad \text{(Hölder's)}$$

Bounding the inner product

Lemma

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}}^2 \leq \eta\langle\nabla_t, \omega_t - \omega_{t+1}\rangle$$

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}}^2 \leq \eta \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\} \langle\nabla_t, \omega_t - \omega_{t+1}\rangle \quad \text{(Combine with Gen. Pinsker's)}$$

$$\frac{1}{4}\|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}} \leq 4\eta \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\}\|\nabla_t\|_{\mathrm{op}} \quad \text{(Hölder's)}$$

Use Hölder's inequality again

$$\langle\nabla_t, \omega_t - \omega_{t+1}\rangle \leq \|\omega_t - \omega_{t+1}\|_{\mathsf{Tr}} \|\nabla_t\|_{\mathrm{op}} \leq 4\eta \max\{\mathrm{Tr}(\omega_t), \mathrm{Tr}(\omega_{t+1})\} \|\nabla_t\|_{\mathrm{op}}^2$$

Summary

- Sublinear regret bound for learning general quantum objects
- Additional matrix analysis results
- Generalized Pinsker's inequality

Summary

- Sublinear regret bound for learning general quantum objects
- Additional matrix analysis results
- Generalized Pinsker's inequality
- Connections to:
    - f-divergence inequalities
    - GPT learning?

Summary

- Sublinear regret bound for learning general quantum objects
- Additional matrix analysis results
- Generalized Pinsker's inequality
- Connections to:
  - f-divergence inequalities
  - GPT learning?
- Future directions:
  - Relax assumptions
  - Connect to shadow tomography
  - Learning infinite-dimensional objects?
  - Other algorithms?

References

Scott Aaronson.
Shadow tomography of quantum states.
Technical Report arXiv:1711.01053, arXiv, November 2018.
arXiv:1711.01053 [quant-ph] type: article.

Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak.
Online learning of quantum states.
In *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018.

Elad Hazan.
Introduction to online convex optimization.
Foundations and Trends® in Optimization, 2(3-4):157–325, 2016.

Yao-Liang Yu.
The Strong Convexity of von Neumann's Entropy.
2013.