

Quantum Online Learning

Defense

Alice Zheng

Virginia Polytechnic Institute and State University

July 6, 2026



1. *Online learning of a panoply of quantum objects*

The rate of information destruction and f -divergence Pinsker inequalities

2. *Sublinear regret is a property of Nature*

GPTs and EJAs

3. *Online unambiguous changepoint detection with unknown quantum states*

4. *Autonomous Hamiltonian certification and changepoint detection*

5. *Local strategies are pretty good at computing Boolean properties of quantum sequences*

Randomness compression in communication networks

Introduction

Quantum computing

- Quantum mechanics: overcomes shortcomings of classical physics
- Gained traction in computer science
 - Better simulation of physics
 - Extend the ability of computers
- Synthesis of theory and implementation
- Finding applications where quantum algorithms have an advantage
- Making better use of existing resources

What insights can be gained by reformulating quantum tasks?

How do we account for the practical limitations of modern quantum computers?

Quantum computing

- Quantum mechanics: overcomes shortcomings of classical physics
- Gained traction in computer science
 - Better simulation of physics
 - Extend the ability of computers
- Synthesis of theory and implementation
- Finding applications where quantum algorithms have an advantage
- Making better use of existing resources

What insights can be gained by reformulating quantum tasks?

How do we account for the practical limitations of modern quantum computers?

Quantum computing

- Quantum mechanics: overcomes shortcomings of classical physics
- Gained traction in computer science
 - Better simulation of physics
 - Extend the ability of computers
- Synthesis of theory and implementation
- Finding applications where quantum algorithms have an advantage
- Making better use of existing resources

What insights can be gained by reformulating quantum tasks?

How do we account for the practical limitations of modern quantum computers?

Interactions between two parties (agent and adversary)

At each round $t \in \{0, \dots, T\}$,

1. Agent chooses action $a_t \in \mathcal{A}$,
2. Adversary picks a loss function f_t ,
3. Agent obtains penalty $f_t(a_t)$ and some information about f_t .

- Choose a_t before f_t is known
- Losses f_t can be different in general
- Losses f_t may be selected **adversarially**

Interactions between two parties (agent and adversary)

At each round $t \in \{0, \dots, T\}$,

1. Agent chooses action $a_t \in \mathcal{A}$,
2. Adversary picks a loss function f_t ,
3. Agent obtains penalty $f_t(a_t)$ and some information about f_t .

- Choose a_t before f_t is known
- Losses f_t can be different in general
- Losses f_t may be selected **adversarially**

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

$$\mathcal{R}_T = \sum_{t=1}^T f_t(a_t) - \min_{a \in \mathcal{A}} \sum_{t=1}^T f_t(a) .$$

- $\mathcal{A}$ and sup are typically clear from context and omitted
- Interested in an upper bound for the worst-case regret
- **Goal:** sublinear regret bound ($\mathcal{R}_T/T \rightarrow 0$ as $T \rightarrow \infty$)

- **Online convex optimization:** set of actions $\mathcal{A}$ and loss functions f_t are convex
- **Bandits:** only $f_t(a_t)$ is revealed
- **Reinforcement learning:** choosing different actions a_t can lead the environment to different states s_t

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

- States can be represented by complex unit vectors $|\psi\rangle \in \mathbb{C}^d$
- Entries are called *amplitudes* – interfere with each other on further operations
- Can evolve state $|\psi\rangle$ by unitary $U \in \mathbb{C}^{d \times d}$ to obtain $U|\psi\rangle$
- Classical states can be probabilistic, described by probability vectors $p \in \mathbb{R}^d$
- To combine amplitudes with probabilities, use *density operators* $\rho \in \mathcal{D}(\mathbb{C}^d)$

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i| \in \mathbb{C}^{d \times d}$$


- Require $\rho \in \text{Pos}(\mathbb{C}^d)$ for nonnegative probabilities, $\text{Tr}(\rho) = 1$ to make them add to 1
- Measurements: $(M_1, \dots, M_k)$ having $M_i \in \text{Pos}(\mathbb{C}^d)$ and $\sum_{i=1}^k M_i = \mathbb{1}$
- Probability to get outcome i when measuring ρ is $\langle M_i, \rho \rangle = \text{Tr}(M_i^* \rho)$

Quantum online learning: concepts


- Online algorithms may be viewed as processing data in chunks
- Quantum data: model via product states / separable states
- Different algorithm structures:



(a) Global



(b) Product



(c) Adaptive

Learning data in a quantum setting

Applying online learning to data from quantum settings

- Recommender systems for quantum data [BLT24]
- Learning descriptions of quantum states [CHL+24]
- Works #1 and #2 herein

[BLT24] Shrigyan Brahmachari, Josep Lumberras, and Marco Tomamichel. “Quantum contextual bandits and recommender systems for quantum data”. *Quantum Machine Intelligence*. 2024

[CHL+24] Xinyi Chen, Elad Hazan, Tongyang Li, Zhou Lu, Xinzhao Wang, and Rui Yang. “Adaptive online learning of quantum states”. *Quantum*. 2024

Online versions of quantum algorithms

Using quantum algorithms for better scaling in online learning settings

- Applied to online learning settings like OCO [HYZL22] and bandits [WYLC21]
- Logarithmic regret in bandits [WZL+23], zero-sum games [GJLW23], reinforcement learning [ZHX+24]

[HYZL22] Jianhao He, Feidiao Yang, Jialin Zhang, and Lvzhou Li. “Quantum algorithm for online convex optimization”. *Quantum Science and Technology*. 2022

[WYLC21] Daochen Wang, Xuchen You, Tongyang Li, and Andrew M. Childs. “Quantum exploration algorithms for multi-armed bandits”. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2021

[WZL+23] Zongqi Wan, Zhijie Zhang, Tongyang Li, Jialin Zhang, and Xiaoming Sun. “Quantum multi-armed bandits and stochastic linear bandits enjoy logarithmic regrets”. *Proceedings of the AAAI Conference on Artificial Intelligence*. 2023

[GJLW23] Minbo Gao, Zhengfeng Ji, Tongyang Li, and Qisheng Wang. “Logarithmic-regret quantum learning algorithms for zero-sum games”. *Advances in Neural Information Processing Systems*. 2023

[ZHX+24] Han Zhong, Jiachen Hu, Yecheng Xue, Tongyang Li, and Liwei Wang. “Provably efficient exploration in quantum reinforcement learning with logarithmic worst-case regret”. *Proceedings of the 41st International Conference on Machine Learning*. 2024

Online versions of quantum algorithms

Using online structure to reduce the footprint of quantum algorithms

- Streaming algorithms [CFL+25]; [GLL+25]
- Works #3, #4, and #5 herein

[CFL+25] Andrew M. Childs, Honghao Fu, Debbie Leung, Zhi Li, Maris Ozols, and Vedang Vyas. “Streaming quantum state purification”. *Quantum*. 2025

[GLL+25] Daniel Grier, Debbie Leung, Zhi Li, Hakop Pashayan, and Luke Schaeffer. “Streaming quantum state purification for general mixed states”. *arXiv:2503.22644*. 2025

Online learning of a panoply of quantum objects

doi:10.1007/s42484-025-00300-5

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

Virginia Tech, UChicago, UIUC

- Mathematical representations of physical objects
- Subsets of positive semidefinite operators

Set	Definition
States	$D(\mathcal{X}) := \{\rho \in \text{Pos}(\mathcal{X}) : \text{Tr}(\rho) = 1\}$
Effects	$\{E \in \text{Pos}(\mathcal{X}) : E \preceq \mathbb{1}\}$
Gram matrices	$\{G \in \text{Pos}(\mathcal{X}) : \forall i, \langle i G i \rangle = 1\}$
Channels	$\{J \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}) : \text{Tr}_{\mathcal{Y}}(J) = \mathbb{1}_{\mathcal{X}}\}$
Int. measurements	$\{R \in \text{Pos}(\mathcal{Y} \otimes \mathcal{X}) : R \preceq \rho \otimes \mathbb{1}_{\mathcal{Y}}, \rho \in D(\mathcal{X})\}$

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

- 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 [Aar18].
 - *Input*: many same quantum d -dimensional *objects* ρ
 - *Input*: m quantum co-objects E_i
 - *Output*: b_i such that $b_i \approx \langle E_i, \rho \rangle$
- Knowing $b_i \rightarrow$ 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: $O(m)$

[Aar18] Scott Aaronson. “Shadow tomography of quantum states”. *Proceedings of the 50th annual ACM SIGACT Symposium on Theory of Computing*. 2018

- **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*.

[Aar18] Scott Aaronson. "Shadow tomography of quantum states". *Proceedings of the 50th annual ACM SIGACT Symposium on Theory of Computing*. 2018

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

Input: $T, \eta > 0$, convex regularizer R , convex and compact set $\mathcal{K} \subseteq \text{Herm}(\mathcal{X})$.

- 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 $f_t : \mathcal{K} \rightarrow \mathbb{R}$ 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**
-

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

Suppose that

- *Loss functions are B -Lipchitz*

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

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

[ACH+18] Scott Aaronson, Xinyi Chen, Elad Hazan, Satyen Kale, and Ashwin Nayak. "Online learning of quantum states". *Advances in Neural Information Processing Systems*. 2018

Algorithm 2 RFTL for Online Learning of Quantum Objects

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

- 1: Set initial hypothesis $\omega_1 \leftarrow \arg \max_{\varphi \in \mathcal{K}} \{S(\varphi)\}$.
- 2: **for** $t \leftarrow 1$ **to** T **do**
- 3: Predict ω_t and incur cost $f_t(\omega_t) := \ell_t(\langle E_t, \omega_t \rangle)$ with $\ell_t : \mathbb{R} \rightarrow \mathbb{R}$ and $E_t \in \mathcal{E}$.
- 4: Let $\ell'_t(x)$ be a sub-derivative of ℓ_t with respect to x and define

$$\nabla_t \leftarrow \ell'_t(\langle E_t, \omega_t \rangle) E_t .$$

- 5: Update decision according to the RFTL rule

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

- 6: **end for**
-

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

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

Lemma

Suppose $\mathcal{K} \in \text{Pos}(\mathcal{X})$ satisfies $\text{Tr}(X) = A$ 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 $\text{Tr}(X) \leq A$, then

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

Objects	Co-objects	Regret upper 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,

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

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

1. Variable traces

- Generalized Pinsker's inequality

$$\text{Tr}(P) \neq \text{Tr}(Q)$$

2. Derivatives of functions over complex operators

- Fréchet differentiation

$$f : \text{Herm}(\mathbb{C}^d) \rightarrow \mathbb{R}$$

1. Variable traces

- Generalized Pinsker's inequality

$$\text{Tr}(P) \neq \text{Tr}(Q)$$

2. Derivatives of functions over complex operators

- Fréchet differentiation

$$f : \text{Herm}(\mathbb{C}^d) \rightarrow \mathbb{R}$$

Theorem

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

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

Theorem

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

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

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

Theorem

For any positive semidefinite P and Q ,

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

- Not as tight as we need...

Theorem

For any positive semidefinite P and Q ,

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

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

Theorem

For any positive semidefinite P and Q ,

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

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

Summary

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

Future directions

- Connections to:
 - f -divergence inequalities [GZB24]
 - GPT learning
- Object-specific bounds
- Efficient scaling in object dimension

[GZB24] Ian George, Alice Zheng, and Akshay Bansal. "Divergence inequalities with applications in ergodic theory". *arXiv:2411.17241*. 2024

Summary

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

Future directions

- Connections to:
 - f -divergence inequalities [GZB24]
 - GPT learning
- Object-specific bounds
- Efficient scaling in object dimension

[GZB24] Ian George, Alice Zheng, and Akshay Bansal. "Divergence inequalities with applications in ergodic theory". *arXiv:2411.17241*. 2024

Sublinear regret is a property of Nature

In preparation


Bharadwaj Sudarsan, Akshay Bansal, Alice Zheng, Jamie Sikora


Virginia Tech

- Sublinear regret for quantum objects $\rightarrow$ “learnable” in some way
- Which theory structures are needed for learnability?
- Further generalize this—**Generalized Probabilistic Theories (GPTs)**
 - Describe sets of states and effects (measurements) producing probabilities
 - General physical theories, possibly beyond quantum
- **States:** S real, finite-dimensional, compact, convex
- **Effects:** for all $e \in \text{Eff}(S)$ and $s \in S$, $e(s) \in [0, 1]$
 - This implies $\text{Eff}(S) \subseteq S^* \cap S^\circ$
- A unique unit effect 1_S satisfies $1_S(s) = 1$ for all $s \in S$



- Sublinear regret for quantum objects $\rightarrow$ “learnable” in some way
- Which theory structures are needed for learnability?
- Further generalize this—**Generalized Probabilistic Theories (GPTs)**
 - Describe sets of states and effects (measurements) producing probabilities
 - General physical theories, possibly beyond quantum
- **States:** S real, finite-dimensional, compact, convex
- **Effects:** for all $e \in \text{Eff}(S)$ and $s \in S$, $e(s) \in [0, 1]$
 - This implies $\text{Eff}(S) \subseteq S^* \cap S^\circ$
- A unique unit effect 1_S satisfies $1_S(s) = 1$ for all $s \in S$





- Does every physical theory permit states to be learnable?
- Is the process of empirical science guaranteed to attain a sufficiently accurate description of reality?


- An *algebra* is a vector space V with a bilinear product $\circ : V \times V \rightarrow V$
- A **Euclidean Jordan Algebra** has an inner product $\langle \cdot, \cdot \rangle$ and identity element $\mathbf{1}$, and
 1. $x \circ y = y \circ x$ and $(x \circ x) \circ (x \circ y) = x \circ ((x \circ x) \circ y)$
 2. $\langle x \circ y, z \rangle = \langle y, x \circ z \rangle$
 3. $\mathbf{1} \circ x = x = x \circ \mathbf{1}$
- Jordan frames $\{c_1, \dots, c_r\}$ satisfy $\sum_{i=1}^r c_i = \mathbf{1}$ and $c_i \circ c_j = 0$ for $i \neq j$
- Jordan frames function analogously to bases; allow taking spectral decompositions and defining Löwner extensions of functions
- One-to-one relationship between EJAs' cones of squares $\{v \in V : \exists x \in V, v = x \circ x\}$ and symmetric cones (Koecher-Vinberg theorem)
- Online learning on symmetric cones is known [CLPV23]

[CLPV23] Ilayda Canyakmaz, Wayne Lin, Georgios Piliouras, and Antonios Varvitsiotis. "Multiplicative updates for online convex optimization over symmetric cones". *arXiv:2307.03136*. 2023

- An *algebra* is a vector space V with a bilinear product $\circ : V \times V \rightarrow V$
- A **Euclidean Jordan Algebra** has an inner product $\langle \cdot, \cdot \rangle$ and identity element $\mathbf{1}$, and
 1. $x \circ y = y \circ x$ and $(x \circ x) \circ (x \circ y) = x \circ ((x \circ x) \circ y)$
 2. $\langle x \circ y, z \rangle = \langle y, x \circ z \rangle$
 3. $\mathbf{1} \circ x = x = x \circ \mathbf{1}$
- Jordan frames $\{c_1, \dots, c_r\}$ satisfy $\sum_{i=1}^r c_i = \mathbf{1}$ and $c_i \circ c_j = 0$ for $i \neq j$
- Jordan frames function analogously to bases; allow taking spectral decompositions and defining Löwner extensions of functions
- One-to-one relationship between EJAs' cones of squares $\{v \in V : \exists x \in V, v = x \circ x\}$ and symmetric cones (Koecher-Vinberg theorem)
- Online learning on symmetric cones is known [CLPV23]

[CLPV23] Ilayda Canakmaz, Wayne Lin, Georgios Piliouras, and Antonios Varvitsiotis. "Multiplicative updates for online convex optimization over symmetric cones". *arXiv:2307.03136*. 2023

- Embed set of states of a GPT into an EJA
- Ψ is an invertible affine transformation
- Begin by proving regret bound when mapping is not necessary
- Bound M , overhead from embedding (via Löwner-John ellipsoids)



- Modification of the symmetric cone MWU algorithm in [CLPV23] via projection

Algorithm 3 Symmetric cone multiplicative weights update with projection

Input: An EJA $(V, \circ, \langle \cdot, \cdot \rangle)$ of rank r , a convex subset S of the trace-one slice of the cone of squares of the EJA, step size $\eta \in (0, 1]$, and cost vectors $-\mathbf{1} \preceq_{\mathcal{K}} m_1, m_2, \dots, m_T \preceq_{\mathcal{K}} \mathbf{1}$.


- 1: $\hat{\omega}_1 \leftarrow \arg \min_{\omega \in \mathcal{K}, \text{Tr}(\omega)=1} \tilde{H}(\omega)$
 - 2: $\omega_1 \leftarrow \arg \min_{\omega \in S} D_{\tilde{H}}(\omega \| \hat{\omega}_1)$
 - 3: **for** $t \in \{1, 2, \dots, T\}$ **do**
 - 4: Predict ω_t and incur cost $f_t(\omega_t) := \ell_t(e_t(\omega_t))$.
 - 5: Suffer loss $\langle m_t, \omega_t \rangle$ and receive cost matrix m_t .
 - 6: $\hat{\omega}_{t+1} \leftarrow \frac{\exp(-\eta \sum_{i=1}^t m_i)}{\text{Tr}(\exp(-\eta \sum_{i=1}^t m_i))}$
 - 7: $\omega_{t+1} \leftarrow \arg \min_{\omega \in S} D_{\tilde{H}}(\omega \| \hat{\omega}_{t+1})$
 - 8: **end for**
-

[CLPV23] Ilayda Canykmaz, Wayne Lin, Georgios Piliouras, and Antonios Varvitsiotis. “Multiplicative updates for online convex optimization over symmetric cones”. *arXiv:2307.03136*. 2023

Theorem

Consider a GPT with its state space S convex, compact and having affine dimension d , as well as an EJA of rank r . There exists an online learning algorithm $\mathcal{A}$ having

$$\text{Regret}_T(\mathcal{A}) = \mathcal{O}\left(LM\sqrt{T\log r}\right), \quad \text{Mistakes}(\mathcal{A}) = \mathcal{O}\left(\frac{L^2d}{\varepsilon^2}\right).$$




$$M_{(X,Y)} := \max_{e \in \text{Eff}(X)} \left\{ \max_{y_1 \in Y} e(y_1) - \min_{y_2 \in Y} e(y_2) \right\}.$$

Theorem

Consider a GPT with its state space S convex, compact and having affine dimension d , as well as an EJA of rank r . There exists an online learning algorithm $\mathcal{A}$ having

$$\text{Regret}_T(\mathcal{A}) = \mathcal{O}\left(LM\sqrt{T\log r}\right), \quad \text{Mistakes}(\mathcal{A}) = \mathcal{O}\left(\frac{L^2d}{\varepsilon^2}\right).$$

$$M_{(X,Y)} := \max_{e \in \text{Eff}(X)} \left\{ \max_{y_1 \in Y} e(y_1) - \min_{y_2 \in Y} e(y_2) \right\}.$$



- **Quantum states**

$$\text{Regret}_T = \mathcal{O}\left(L\sqrt{T\log(d)}\right) , \quad \text{Mistakes} = \mathcal{O}(L^2\log(d)/\varepsilon^2) .$$

- **Quantum channels**

$$\text{Regret}_T = \mathcal{O}(Ld_Ad_B\sqrt{T}) , \quad \text{Mistakes} = \mathcal{O}(L^2d_A^2d_B^2/\varepsilon^2) .$$

- **Quantum effects**

$$\text{Regret}_T = \mathcal{O}(Ld\sqrt{T}) , \quad \text{Mistakes} = \mathcal{O}(L^2d^2/\varepsilon^2) .$$

- **Elliptical GPTs**

$$\text{Regret}_T = \mathcal{O}(L\sqrt{T}) , \quad \text{Mistakes} = \mathcal{O}(L^2/\varepsilon^2) .$$

Summary

- Online learning of GPT states
- Leverage known MWU algorithm for symmetric cones
- Mappings between GPTs and EJAs
- Sublinear regret as a universal concept
- Reinforce validity of GPTs as physical theories

Future directions

- Choices other than $r = 2$
- Object-specific bounds to utilize additional structure

Summary

- Online learning of GPT states
- Leverage known MWU algorithm for symmetric cones
- Mappings between GPTs and EJAs
- Sublinear regret as a universal concept
- Reinforce validity of GPTs as physical theories

Future directions

- Choices other than $r = 2$
- Object-specific bounds to utilize additional structure


Online unambiguous changepoint detection with unknown quantum states

In preparation

Alice Zheng, Jamie Sikora, Sarvagya Upadhyay

Virginia Tech, Fujitsu Research of America

- Detecting abrupt changes in data
- Used classically for: intrusion detection, quality control
- Quantum setup:



- Known/**unknown** states, identification / **detection**, minimal-error / **unambiguous**

- Parameters: states $|A\rangle$ and ρ_B , length n
 - Described by probability measures
 - No false alarms: unambiguous discrimination
 - Maximize (average) detection probability
- **Bottom-up**: simple algorithms
 - Increase detection probability
 - **Top-down**: optimal algorithm
 - Simplify constructions

- Parameters: states $|A\rangle$ and ρ_B , length n
- Described by probability measures
- No false alarms: unambiguous discrimination
- Maximize (average) detection probability
- **Bottom-up**: simple algorithms
 - Increase detection probability
- **Top-down**: optimal algorithm
 - Simplify constructions

- Corresponds to POVM elements

$$\frac{1}{2}\mathbb{1} + \frac{(-1)^i}{2}W$$



- Measurement probabilities given by

$$p^{(i)} = \frac{1}{2}(1 \mp \langle \rho_A, \rho_B \rangle)$$

- Post-measurement state is

$$\rho^{(i)} = \frac{1}{p^{(i)}}(\langle i| \otimes \mathbb{1})\rho(|i\rangle \otimes \mathbb{1})$$

- Detecting *changes* between states
- Detecting *overlap* between states



- Keep one state in memory
- SWAP test incoming state with state in memory
- When given i copies of the changed state:

$$p_{\text{succ}} \geq \frac{2(1 - 4^{-i})}{3}(1 - \langle \rho_A, \rho_B \rangle)$$

$$n_{\text{delay}} = \frac{4^{i+1} - 3i - 4}{3(4^i - 1)}$$

- Success probability tends to $2/3$
- Delay tends to $4/3$

- Keep more states in memory
 - SWAP test incoming with one from memory
 - SWAP test the two states above with two from memory
 - And so on...
- Better performance with larger memory k

$$p_{\text{succ}} = (1 - 2^{-k})(1 - \langle \rho_A, \rho_B \rangle)$$

- Orthogonal states: approaches perfect detection

- Semidefinite program with solution independent of changepoint details
- Solution $X = \Pi_{\text{null}(A)} = \mathbb{1} - \Pi_{\mathcal{X}^{\otimes n}}$ can be used for other problems relating to the symmetric subspace

$$\begin{aligned}
 B &= \int \rho_A^{\otimes \nu} \otimes \rho_B^{\otimes n-\nu} d\lambda(\nu, \rho_A, \rho_B) & \text{maximize: } \langle B, X \rangle \\
 A &= \int \rho_A^{\otimes n} d\mu(\rho_A), & \text{subject to: } \langle A, X \rangle = 0 \\
 & & 0 \preceq X \preceq \mathbb{1},
 \end{aligned}$$

- Success probability analytically given by

$$1 - \sum_{x=0}^{\lfloor n/2 \rfloor} \sum_{\nu=x}^{n-x} p(\nu) \binom{n}{\nu}^{-1} \binom{\nu}{x} \binom{n-\nu}{x} |\langle \psi_A, \psi_B \rangle|^{2x},$$


- Semidefinite program with solution independent of changepoint details
- Solution $X = \Pi_{\text{null}(A)} = \mathbb{1} - \Pi_{\mathcal{X}^{\otimes n}}$ can be used for other problems relating to the symmetric subspace

$$\begin{aligned}
 B &= \int \rho_A^{\otimes \nu} \otimes \rho_B^{\otimes n-\nu} d\lambda(\nu, \rho_A, \rho_B) & \text{maximize: } \langle B, X \rangle \\
 A &= \int \rho_A^{\otimes n} d\mu(\rho_A), & \text{subject to: } \langle A, X \rangle = 0 \\
 & & 0 \preceq X \preceq \mathbb{1},
 \end{aligned}$$

- Success probability analytically given by


$$1 - \sum_{x=0}^{\lfloor n/2 \rfloor} \sum_{\nu=x}^{n-x} p(\nu) \binom{n}{\nu}^{-1} \binom{\nu}{x} \binom{n-\nu}{x} |\langle \psi_A, \psi_B \rangle|^{2x},$$

- **Basic** close to $2/3$ most of the time
- **Extended basic** and **optimal** behave similarly, close performance
- Can interpolate between **basic** and **extended basic** by increasing amount of quantum memory




- Sum of permutations

$$\Pi_{\mathcal{X}^{\otimes n}} = \sum_{k=1}^{n!} W_k$$




- Create uniform superposition
- Control each permutation
- Recombine to obtain sum

Ancilla	$\mathcal{O}(n \log n)$
CSWAPs	$n \times \mathcal{O}(n!)$
Depth	$n \times \mathcal{O}(n!)$
Depth*	$n \times \mathcal{O}(n!)$




- Circuit is recursive – can split into steps
- May be viewed as an “online” circuit (with growing memory)
- Can recycle the ancilla between steps

Ancilla	$\mathcal{O}(\log n)$
CSWAPs	$n \times \mathcal{O}(n^2 \log n)$
Depth	$n \times \mathcal{O}(n^2 \log n)$
Depth*	$n \times \mathcal{O}(n \log n)$




- Circuit is recursive – can split into steps
- May be viewed as an “online” circuit (with growing memory)
- Can recycle the ancilla between steps


Ancilla	$\mathcal{O}(\log n)$
CSWAPs	$n \times \mathcal{O}(n^2 \log n)$
Depth	$n \times \mathcal{O}(n^2 \log n)$
Depth*	$n \times \mathcal{O}(n \log n)$



- SWAP all memory registers with the new one
- Note that the order of registers is rearranged between steps


Ancilla	$\mathcal{O}(\log n)$
CSWAPs	$n \times \mathcal{O}(n \log n)$
Depth	$n \times \mathcal{O}(n \log n)$
Depth*	$n \times \mathcal{O}(n \log n)$





- Decompose ancilla into qubits
- Control on individual qubits
- Account for identity permutation

Ancilla	$\mathcal{O}(\log n)$
CSWAPs	$n \times \mathcal{O}(n)$
Depth	$n \times \mathcal{O}(n)$
Depth*	$n \times \mathcal{O}(\log n)$





- Decompose ancilla into qubits
- Control on individual qubits
- Account for identity permutation

Ancilla	$\mathcal{O}(\log n)$
CSWAPs	$n \times \mathcal{O}(n)$
Depth	$n \times \mathcal{O}(n)$
Depth*	$n \times \mathcal{O}(\log n)$

Theorem (Informal)

When their size is a power of 2, the basic and optimal circuits act equivalently on quantum memory in the symmetric subspace.



- Happens in several restricted cases:
 1. Instantaneous detection probability, and changepoint location is a power of 2
 2. Memory limit in both algorithms that is a power of 2
- Note: not known whether memory-limited optimal circuit is still optimal



Theorem (Informal)

When their size is a power of 2, the basic and optimal circuits act equivalently on quantum memory in the symmetric subspace.

- Happens in several restricted cases:
 1. Instantaneous detection probability, and changepoint location is a power of 2
 2. **Memory limit** in both algorithms that is a power of 2
- *Note:* not known whether memory-limited optimal circuit is still optimal



Summary

- Simple circuits with memory-vs-performance tradeoff
- Optimal circuits independent of changepoint details
- Equivalence in specific cases
- Connection of the no-false-alarm constraint to the symmetric subspace

Future directions

- Changepoint in unitaries/channels
- No known priors
- Further memory reduction
- Memory-limited optimal still optimal?

Summary

- Simple circuits with memory-vs-performance tradeoff
- Optimal circuits independent of changepoint details
- Equivalence in specific cases
- Connection of the no-false-alarm constraint to the symmetric subspace

Future directions

- Changepoint in unitaries/channels
- No known priors
- Further memory reduction
- Memory-limited optimal still optimal?

Autonomous Hamiltonian certification and changepoint detection

arXiv:2603.26655

Steven T. Flammia, Dmitrii Khitrin, Muzhou Ma, Jamie Sikora,
Yu Tong, Alice Zheng

Virginia Tech, Phasecraft, Duke, Caltech

- A system that evolves over time can be described by the Schrödinger equation:

$$i\hbar \frac{d}{dt} |\psi(t)\rangle = \mathbf{H} |\psi(t)\rangle$$

- $\mathbf{H}$ is a **Hamiltonian**: Hermitian linear operator
- Described by unitary evolution when H is time-independent

$$|\psi(t)\rangle = e^{-iHt/\hbar} |\psi(0)\rangle =: U(t) |\psi(0)\rangle$$



- Scenarios requiring modeling quantum interactions:
 - Quantum chemistry
 - Material design
 - Many-body physics
 - etc.
- A promising application is to simulate these on a quantum computer
- Dynamics of quantum devices can likewise be described by Hamiltonians

- A device is promised to implement some *ideal* Hamiltonian H_0 (known)
- Instead, the implemented Hamiltonian is H
- Need to distinguish between the following cases:

(Near): $\|H - H_0\|_F \leq \epsilon_1$,

(Far): $\|H - H_0\|_F \geq \epsilon_2$

- $\|\cdot\|_F$ is the *normalized*^a Frobenius norm
- Design an algorithm that with high probability:
 - **Accepts** in the (Near) case,
 - **Rejects** in the (Far) case
- Minimize samples and total evolution time
- Applicable to benchmarking / validation





^aDivided by $\sqrt{d}$, which accounts for the scaling of the Frobenius norm under stochastic noise on every parameter.

- A device is promised to implement some *ideal* Hamiltonian H_0 (known)
- Instead, the implemented Hamiltonian is H
- Need to distinguish between the following cases:

(Near): $\|H - H_0\|_F \leq \epsilon_1$,

(Far): $\|H - H_0\|_F \geq \epsilon_2$

- $\|\cdot\|_F$ is the *normalized*^a Frobenius norm
- Design an algorithm that with high probability:
 - **Accepts** in the (Near) case,
 - **Rejects** in the (Far) case
- Minimize samples and total evolution time
- Applicable to benchmarking / validation



^aDivided by $\sqrt{d}$, which accounts for the scaling of the Frobenius norm under stochastic noise on every parameter.

- How does a quantum device certify itself?
- *Autonomous*: no external assistance
- *Single-qubit* state preparation and measurements
- Examples of resources this prohibits:
 - Second device / trusted ancilla,
 - Evolution $e^{-iH_0 t}$ by the ideal Hamiltonian,
 - Entangled state preparation and measurements,
 - Controlled / inverse-time evolution,
 - Trotterization
- Hamiltonians have *operator norm bounded by M*
 - n -qubit, traceless, constant bound on coefficients

- Algorithm for certifying quantum *states*
 1. Randomly sample $k \in \{1, \dots, n\}$,
 2. Measure qubits 1 to $(k - 1)$ in computational basis,
 3. Measure qubits $(k + 1)$ to n in *decision tree* basis,
 4. Certify on qubit k
- Only uses single-qubit gates, no controlled / inverse-time evolution
- $\mathcal{O}(n\xi^{-1} \ln(1/\delta))$ samples for precision $\xi_1 = \xi/(2n)$, $\xi_2 = \xi$
- Dependence on system size n not necessary with multi-qubit gates
 - n present in sample complexity and threshold ratio
 - Future direction (out of scope)

[GHO25] Meghal Gupta, William He, and Ryan O'Donnell. "Few single-qubit measurements suffice to certify any quantum state". *arXiv:2506.11355*. 2025

- Algorithm for certifying quantum *states*
 1. Randomly sample $k \in \{1, \dots, n\}$,
 2. Measure qubits 1 to $(k - 1)$ in computational basis,
 3. Measure qubits $(k + 1)$ to n in *decision tree* basis,
 4. Certify on qubit k
- Only uses single-qubit gates, no controlled / inverse-time evolution
- $\mathcal{O}(n\xi^{-1} \ln(1/\delta))$ samples for precision $\xi_1 = \xi/(2n)$, $\xi_2 = \xi$
- Dependence on system size n not necessary with multi-qubit gates
 - n present in sample complexity and threshold ratio
 - Future direction (out of scope)

[GHO25] Meghal Gupta, William He, and Ryan O'Donnell. "Few single-qubit measurements suffice to certify any quantum state". *arXiv:2506.11355*. 2025

- More total parameters (compared to pure states)
- More flexibility: choose input states and measurements
 1. Sample stabilizer state $|\psi_0\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n}$,
 2. Evolve by H for time $t = cM^{-1}$,
 3. Certify using [GHO25] w.r.t. a classically-computed state,
 4. Repeat N times and perform a likelihood ratio test
- Relating distance $\|H_0 - H\|_F$ and infidelity $1 - F(\rho, \sigma)$
 - Accumulated differences described by evolution $W(t) = e^{iH_0 t} e^{-iHt}$,
 - Bound $\mathbb{1}$ term in Pauli decomposition,
 - Analyze 4th order Taylor expansion of $|\text{Tr}[W(t)]|^2$,
 - Proof uses normal ordering

[GHO25] Meghal Gupta, William He, and Ryan O'Donnell. "Few single-qubit measurements suffice to certify any quantum state". *arXiv:2506.11355*. 2025

- More total parameters (compared to pure states)
- More flexibility: choose input states and measurements
 1. Sample stabilizer state $|\psi_0\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle, |i\rangle, |-i\rangle\}^{\otimes n}$,
 2. Evolve by H for time $t = cM^{-1}$,
 3. Certify using [GHO25] w.r.t. a classically-computed state,
 4. Repeat N times and perform a likelihood ratio test
- Relating distance $\|H_0 - H\|_F$ and infidelity $1 - F(\rho, \sigma)$
 - Accumulated differences described by evolution $W(t) = e^{iH_0 t} e^{-iHt}$,
 - Bound $\mathbb{1}$ term in Pauli decomposition,
 - Analyze 4th order Taylor expansion of $|\text{Tr}[W(t)]|^2$,
 - Proof uses normal ordering

[GHO25] Meghal Gupta, William He, and Ryan O'Donnell. "Few single-qubit measurements suffice to certify any quantum state". *arXiv:2506.11355*. 2025



- Total sample complexity and evolution time

$$N = \mathcal{O}(nM^2\epsilon^{-2} \ln(1/\delta)) , \quad T = \mathcal{O}(nM\epsilon^{-2} \ln(1/\delta))$$


- Tolerance thresholds $\epsilon_1 = \mathcal{O}(\epsilon/\sqrt{n})$, $\epsilon_2 = \epsilon$
- Classical computation polynomial in the number of nonzero Pauli terms $\mathbf{m}$
- Comparison to prior work [GJW+26]:
 - Linear in n ; linear M factor instead of $\mathbf{m}^{3/2}$
 - n dependence removeable via [FL11] (at the cost of worse ϵ -scaling)

[GJW+26] Minbo Gao, Zhengfeng Ji, Qisheng Wang, Wenjun Yu, and Qi Zhao. "Quantum Hamiltonian certification". *Proceedings of the 2026 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 2026



[FL11] Steven T. Flammia and Yi-Kai Liu. "Direct fidelity estimation from few Pauli measurements". *Physical Review Letters*. 2011




- Hamiltonian $H(t)$ can in general vary with time
- Analysis becomes more complicated
- **Quantize**: sequence of Hamiltonians $H_1, H_2, \dots$
- **Frame challenge**: does this lose information?
 - High-frequency components missed
 - Bound spectrum of the Hamiltonian



- A change could happen at any sequence index ν
- Need to distinguish between the following cases:
 (No change): $\|H_i - H_0\|_F \leq \epsilon_1$ for all $i \in \{1, 2, \dots\}$,
 (Change): $\|H_i - H_0\|_F \geq \epsilon_2$ for all
 $i \in \{\nu + 1, \nu + 2, \dots\}$
- Design an algorithm that with high probability:
 - **Keeps running** in the (No change) case,
 - **Terminates rapidly** in the (Change) case
- Minimize detection delay $\bar{\mathbb{E}}$;
- Maximize runtime to false alarm $\mathbb{E}_\infty$







- Noise / anomaly / threat detection in an *online* fashion
- **Recalibration** problem:
 - Changes in $H(t)$ represent accumulated miscalibrations,
 - An (expensive) recalibration procedure can be run to reset $H(t)$ to H_0 ,
 - When to trigger recalibration?

- No assumptions on the stochasticity of the miscalibrations
 - Loss to an environment,
 - Collected by an eavesdropper,
 - Manipulated by an **adversary**
- Adversary could have full information about the used architecture / protocol
- Models various types of attacks
 - Extraction of sensitive information,
 - Manipulation of outcomes,
 - Denial of service


- Classical procedure for continuous (online) inspection
 - Monitor for changes / anomalies in time series data
 - Terminate rapidly following a sudden change
- Initialize a cumulative score $S_0 = 0$. At step i :
 1. Receive time-series data X_i ,
 2. Compute log-likelihood ratio $Z_i = \ln \frac{dQ}{dP}(X_i)$,
 3. Update $S_i = \max\{0, S_{i-1} + Z_i\}$,
 4. Terminate if $S_i \geq h$
- Score stays near zero pre-change;
- Increases linearly post-change



- Equivalent to a parallel series of open-ended ratio tests:
 - Terminate when change is e^h times more likely
 - Predicts MLE for past change location
- Known optimal scaling for delay

$$\mathbb{E}[\max\{0, N - \nu\}] \geq \frac{\ln \mathbb{E}_\infty[N]}{D(Q||P)} \cdot (1 + o(1))$$

- **Robust** over classes of distributions
 - Analysis via least-favorable distributions [MF19]



[MF19] Timothy L. Molloy and Jason J. Ford. “Minimax robust quickest change detection in systems and signals with unknown transients”. *IEEE Transactions on Automatic Control*. 2019


[Lor71] Gary Lorden. “Procedures for reacting to a change in distribution”. *The annals of mathematical statistics*. 1971

- Use $P = \text{Ber}(\xi/2n)$, $Q = \text{Ber}(\xi/n)$
- Relation between detection delay and false alarm runtime

$$\mathbb{E}[\max\{0, N - \nu\}] \leq \mathcal{O}\left(\frac{nM^2 \ln \mathbb{E}_\infty[N]}{\epsilon^2}\right),$$

$$\mathbb{E}[\max\{0, T - t\nu\}] \leq \mathcal{O}\left(\frac{nM \ln(M\mathbb{E}_\infty[T])}{\epsilon^2}\right).$$

- Tolerance thresholds $\epsilon_1 = \mathcal{O}(\epsilon/\sqrt{n})$, $\epsilon_2 = \epsilon$



Summary

- *Autonomous* miscalibration / anomaly / threat detection
- **Certification** for offline settings
- Online **changepoint** detection
- Quantum algorithm sensitivity linear in operator norm M
- Classical simulation complexity linear in sparsity m

Future directions

- Dependence on the number of qubits n (can be avoided via [FL11])
- Analysis of spectrum bounds for time-dependent Hamiltonians
- Integration of changepoint methods with quantum data
- Adaptive choice of input states
- Open quantum systems and Lindbladian mechanics

[FL11] Steven T. Flammia and Yi-Kai Liu. "Direct fidelity estimation from few Pauli measurements". *Physical Review Letters*. 2011

Summary

- *Autonomous* miscalibration / anomaly / threat detection
- **Certification** for offline settings
- Online **changepoint** detection
- Quantum algorithm sensitivity linear in operator norm M
- Classical simulation complexity linear in sparsity m

Future directions

- Dependence on the number of qubits n (can be avoided via [FL11])
- Analysis of spectrum bounds for time-dependent Hamiltonians
- Integration of changepoint methods with quantum data
- Adaptive choice of input states
- Open quantum systems and Lindbladian mechanics


[FL11] Steven T. Flammia and Yi-Kai Liu. “Direct fidelity estimation from few Pauli measurements”. *Physical Review Letters*. 2011

Local strategies are pretty good at computing Boolean properties of quantum sequences

arXiv:2603.05452

Tathagata Gupta, Ankith Mohan, Shayee Murshid, Vincent Russo,
Jamie Sikora, Alice Zheng

IIT Madras, Virginia Tech, ISI Kolkata, Unitary Foundation



- **Minimum-error** or unambiguous
- Uniform priors
- Differences between local, adaptive local (online), and global strategies
- Use cases:
 - Quantum reading of classical memories [Pir11]
 - Distributed sensing
 - Quantum-enhanced pattern recognition
 - Bounded-storage cryptography
- Performance guarantee in the memoryless case?
- For which functions is the memoryless strategy optimal?

[Pir11] Stefano Pirandola. “Quantum reading of a classical digital memory”. *Physical Review Letters*. 2011

- Discrimination between the states

$$\rho_0 = \frac{1}{|S_0|} \sum_{x \in S_0} |\psi_x\rangle\langle\psi_x| \quad \text{and} \quad \rho_1 = \frac{1}{|S_1|} \sum_{x \in S_1} |\psi_x\rangle\langle\psi_x|$$

- PGM measurement operators are obtained as follows:

$$M_i^{\text{PGM}} = \rho_{\text{avg}}^{-1/2} (p_i \rho_i) \rho_{\text{avg}}^{-1/2}, \quad \rho_{\text{avg}} := \sum_{i=1}^n p_i \rho_i$$

- Performance guarantee by Barnum and Knill's theorem [BK02]

[BK02] Howard Barnum and Emanuel Knill. "Reversing quantum dynamics with near-optimal quantum and classical fidelity". *Journal of Mathematical Physics*. 2002

- Decompose the average state

$$\rho_{\text{avg}} = p_0 \rho_0 + p_1 \rho_1 = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} |\psi_x\rangle \langle \psi_x| = \left(\frac{1}{2} |\psi_0\rangle \langle \psi_0| + \frac{1}{2} |\psi_1\rangle \langle \psi_1| \right)^{\otimes n} = \rho^{\otimes n}$$

- Decompose the measurement operators, meaning it is memoryless

$$M_{\sigma_i}^{\text{PGM}} = \sum_{y_1 \dots y_n \in S_i} M_{y_1} \otimes \dots \otimes M_{y_n} = \sum_{y \in S_i} M_y, \quad M_i = \frac{1}{2} \rho^{-1/2} |\psi_i\rangle \langle \psi_i| \rho^{-1/2}$$

Theorem

For all Boolean functions f , positive integers n , and $s \in (0,1)$, the performance of the greedy strategy can be bounded below by the following relation

$$P(n, f, s, \text{global})^2 \leq P(n, f, s, \text{greedy}) .$$

- Decompose the average state

$$\rho_{\text{avg}} = p_0 \rho_0 + p_1 \rho_1 = \frac{1}{2^n} \sum_{x \in \{0,1\}^n} |\psi_x\rangle \langle \psi_x| = \left(\frac{1}{2} |\psi_0\rangle \langle \psi_0| + \frac{1}{2} |\psi_1\rangle \langle \psi_1| \right)^{\otimes n} = \rho^{\otimes n}$$

- Decompose the measurement operators, meaning it is memoryless

$$M_{\sigma_i}^{\text{PGM}} = \sum_{y_1 \dots y_n \in S_i} M_{y_1} \otimes \dots \otimes M_{y_n} = \sum_{y \in S_i} M_y, \quad M_i = \frac{1}{2} \rho^{-1/2} |\psi_i\rangle \langle \psi_i| \rho^{-1/2}$$

Theorem

For all Boolean functions f , positive integers n , and $s \in (0, 1)$, the performance of the greedy strategy can be bounded below by the following relation

$$P(n, f, s, \text{global})^2 \leq P(n, f, s, \text{greedy}) .$$

Theorem

The greedy strategy is globally optimal, that is

$$P(n, f, s, \text{greedy}) = P(n, f, s, \text{global})$$

for all $s \in (0, 1)$, if and only if f is affine, that is

$$f(x_1, \dots, x_n) = b_0 \oplus b_1 x_1 \oplus b_2 x_2 \oplus \dots \oplus b_n x_n .$$

- Example: $f = \text{XOR}$ is affine, therefore local same as global
- For all other f the success probabilities differ

- Sufficiency (affine $\Rightarrow$ optimal):
 - Same decomposition allows computing matching success probabilities
 - Global bias = local bias exponentiated by number of considered bits
- Necessity (optimal $\Rightarrow$ affine):
 - PGM optimality equivalent to $[G_{X'B'}, \tau_{X'B'}] = 0$, where $G_{X'B'}$ is the generalized Gram matrix and $\tau_{X'B'}$ its block-diagonal [IRS16]
 - Writing $G = \frac{1}{2^n} \begin{pmatrix} \Gamma_0 & \Gamma' \\ \Gamma'^* & \Gamma_1 \end{pmatrix}$, optimality implies $\Gamma_0 \Gamma' = \Gamma' \Gamma_1$
 - In the hypercube graph, implies a counting condition

$$|\{u \in S_0 : d(x, u) + d(u, y) = L\}| = |\{v \in S_1 : d(x, v) + d(v, y) = L\}|$$

- Hypercube partitions related by a coordinate flip, implying affine structure

[IRS16] Raban Iten, Joseph M. Renes, and David Sutter. “Pretty good measures in quantum information theory”. *IEEE Transactions on Information Theory*. 2016


- Sufficiency (affine $\Rightarrow$ optimal):
 - Same decomposition allows computing matching success probabilities
 - Global bias = local bias exponentiated by number of considered bits
- Necessity (optimal $\Rightarrow$ affine):
 - PGM optimality equivalent to $[G_{X'B'}, \tau_{X'B'}] = 0$, where $G_{X'B'}$ is the generalized Gram matrix and $\tau_{X'B'}$ its block-diagonal [IRS16]
 - Writing $G = \frac{1}{2^n} \begin{pmatrix} \Gamma_0 & \Gamma' \\ \Gamma'^* & \Gamma_1 \end{pmatrix}$, optimality implies $\Gamma_0 \Gamma' = \Gamma' \Gamma_1$
 - In the hypercube graph, implies a counting condition

$$|\{u \in S_0 : d(x, u) + d(u, y) = L\}| = |\{v \in S_1 : d(x, v) + d(v, y) = L\}|$$

- Hypercube partitions related by a coordinate flip, implying affine structure


[IRS16] Raban Iten, Joseph M. Renes, and David Sutter. “Pretty good measures in quantum information theory”. *IEEE Transactions on Information Theory*. 2016

- Highly unbalanced
- As $n \rightarrow \infty$, both strategies approach perfect success
- Gap larger for states with high overlap
- Gap becomes smaller as n increases



- Balanced function
- Largest gap for medium overlap values
- Gap increases with n
- Greedy strategy asymptotically approaches

$$\frac{1}{2} + \frac{1}{\pi} \arcsin(\sqrt{1-s^2})$$



Summary

- Local vs. global strategies for (minimum-error) evaluating Boolean functions on a sequence of quantum states
- Local greedy strategy performs as well as the PGM
- Greedy is optimal exactly for affine functions

Future directions

- Similar guarantees for unambiguous evaluation?
 - Adaptive unambiguous algorithm found via dynamic programming
- Adaptive local strategies
- Online variant of the setting

Summary

- Local vs. global strategies for (minimum-error) evaluating Boolean functions on a sequence of quantum states
- Local greedy strategy performs as well as the PGM
- Greedy is optimal exactly for affine functions

Future directions

- Similar guarantees for unambiguous evaluation?
 - Adaptive unambiguous algorithm found via dynamic programming
- Adaptive local strategies
- Online variant of the setting

Concluding remarks

- General learnability feasible; more efficient through object-specific bounds?
- Tradeoffs between allowed memory size (classical or quantum) and algorithm performance
 - Performance increases may become marginal
 - In settings with specific structure, additional memory does not help
- Remaining questions related to memory constraints:
 - Do memory limits affect optimality?
 - How to reduce amount of memory used?

Comparison to known lower bounds

Task	Metric	Upper bound	Lower bound	Source
State learning	Regret	$\mathcal{O}(\sqrt{nT})$	$\Omega(\sqrt{nT})$	[AHK12]
	Mistakes	$\mathcal{O}(n/\epsilon^2)$	$\Omega(n/\epsilon^2)$	[Aar07]
Channel learning	Regret	$\mathcal{O}(2^n \sqrt{nT}), \mathcal{O}(4^n \sqrt{T})$	—	—
	Mistakes	$\mathcal{O}(n 4^n / \epsilon^2), \mathcal{O}(16^n / \epsilon^2)$	$\Omega(2^n)$	[RCEK24]
Hamiltonian chpt. (fixed strategy)	Delay	$\mathcal{O}\left(\frac{n M^2 \ln \mathbb{E}_\infty[N]}{\epsilon^2}\right)$	$\Omega\left(\frac{M^2 \ln \mathbb{E}_\infty[N]}{\epsilon^2 \ln n}\right)$	[FHC23]

[AHK12] Sanjeev Arora, Elad Hazan, and Satyen Kale. “The multiplicative weights update method: a meta-algorithm and applications”. *Theory of Computing*. 2012

[Aar07] Scott Aaronson. “The learnability of quantum states”. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*. 2007

[RCEK24] Asad Raza, Matthias C. Caro, Jens Eisert, and Sumeet Khatri. “Online learning of quantum processes”. *arXiv:2406.04250*. 2024

[FHC23] Marco Fanizza, Christoph Hirche, and John Calsamiglia. “Ultimate limits for quickest quantum change-point detection”. *Physical Review Letters*. 2023

- *Object learning*: which objects are learnable with regret polylogarithmic in dimension?
- *Changepoint*: increasing efficiency, open quantum systems
- *Sequences*: accounting for different priors, unambiguous discrimination
- A general framework for optimization / algorithm design with limited quantum memory?
 - Analogs for adaptive algorithms [OZN+26]

[OZN+26] Ties-Albrecht Ohst, Shijun Zhang, Hai Chau Nguyen, Martin Plávala, and Marco Túlio Quintino. “Characterising memory in quantum channel discrimination via constrained separability problems”. *Quantum*. 2026

Thank you!

See thesis for more details and a full list of references