# Quantum Changepoint Detection:

## A Statistical Framework for Cyber-Physical Threat Recognition

Alice Zheng (on behalf of Steven T. Flammia)

Virginia Polytechnic Institute and State University

April 28, 2026

# Autonomous Hamiltonian certification and changepoint detection

arXiv:2603.26655

---

Steven T. Flammia[1,2], Dmitrii Khitrin[3], Muzhou Ma[4], Jamie Sikora[1], Yu Tong[3], and Alice Zheng[1]

1 VIRGINIA TECH

2 PHASECRAFT

3 Duke UNIVERSITY

4 Caltech

# Hamiltonians

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

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

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

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

# Hamiltonians: Use cases

- 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

# Hamiltonian Certification

# Certification: Setting

- 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:

$$\begin{aligned} \text{(Near):} \quad & \|H - H_0\|_F \le \epsilon_1, \\ \text{(Far):} \quad & \|H - H_0\|_F \ge \epsilon_2 \end{aligned}$$

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

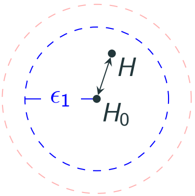

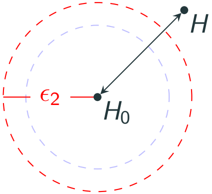

---

[a] Divided by $\sqrt{d}$, which accounts for the scaling of the Frobenius norm under stochastic noise on every parameter.

# Certification: Assumptions and practicality

- 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

# Quantum state certification [GHO25]

- Algorithm for certifying quantum *states*
  1. Randomly sample $k \in \{1, \ldots, n\}$,
  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*. 2025

# Hamiltonian certification procedure

- 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 $I$ term in Pauli decomposition,
    - Analyze 4th order Taylor expansion of $|\mathrm{tr}[W(t)]|^2$,
    - Proof uses normal ordering

---

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

# Certification: Results

- 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 **m**
- Comparison to prior work [Gao+26]:
  - Linear in $n$; linear $M$ factor instead of $\mathbf{m}^{3/2}$
  - $n$ dependence removeable via [FL11] (at the cost of worse $\epsilon$-scaling)

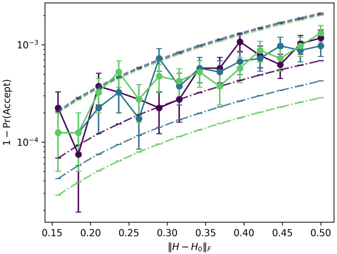

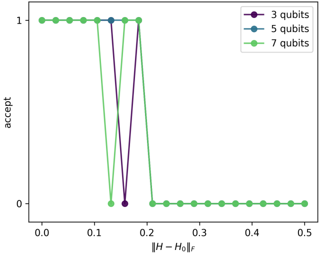

---

[Gao+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 Changepoint

# Time-dependent Hamiltonians / Quantization

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

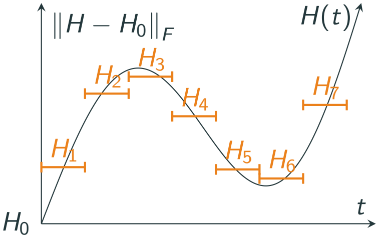

# Changepoint: Setting

- A change could happen at any sequence index $\nu$
- Need to distinguish between the following cases:

$$\begin{aligned} &\text{(No change):} && \|H_i - H_0\|_F \le \epsilon_1 \text{ for all } i \in \{1, 2, \ldots\}, \\ &\text{(Change):} && \|H_i - H_0\|_F \ge \epsilon_2 \text{ for all } \\ & && i \in \{\nu + 1, \nu + 2, \ldots\} \end{aligned}$$

- Design an algorithm that with high probability:
  - **Keeps running** in the (No change) case,
  - **Terminates** *rapidly* in the (Change) case
- Minimize detection delay $\overline{\mathbb{E}}$;
- Maximize runtime to false alarm $\mathbb{E}_\infty$

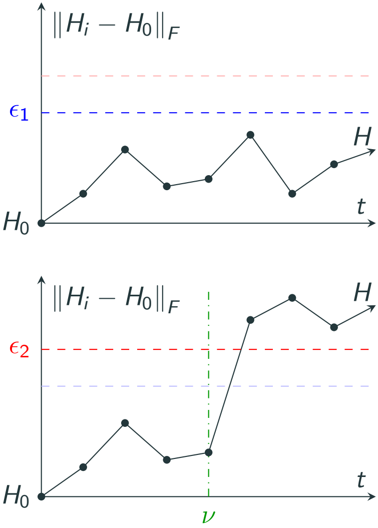

# Changepoint: Applications

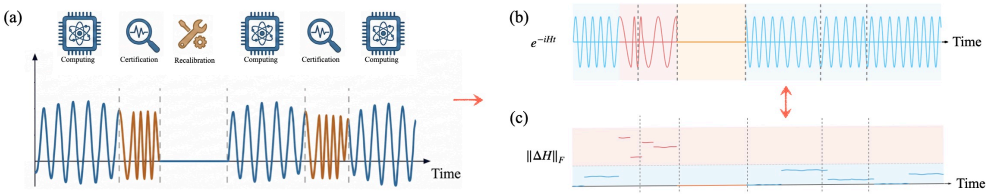

- Noise / anomaly / threat detection in an *online* fashion
- Recalibration problem:
  - Changes in $H(t)$ represent accumulated miscalibrations,
  - An (expensive) recalibration procedure can be run to reset $H(t)$ to $H_0$,
  - When to trigger recalibration?

# Adversarial framework

- 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

# CUSUM [Pag54]

- 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

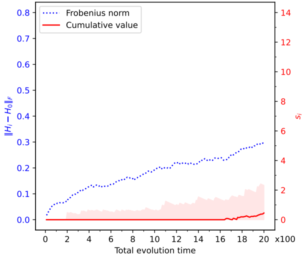

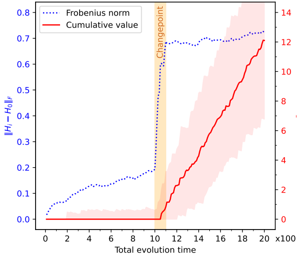

---

[Pag54] Ewan S. Page. “Continuous inspection schemes”. *Biometrika*. 1954

# CUSUM: Optimal scaling [Lor71]

- 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

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

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

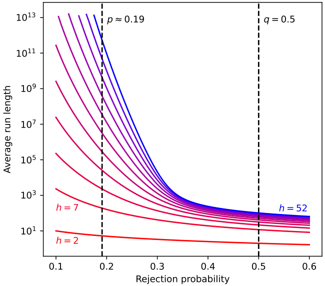

---

[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

# Changepoint: Results

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

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

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

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

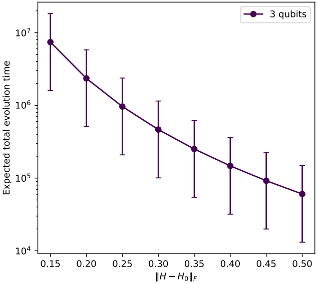

# Conclusion / Future directions

- *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 $\mathbf{m}$

- Dependence on the number of qubits $n$
- Analysis of spectrum bounds for time-dependent Hamiltonians
- Integration of changepoint methods with quantum data
- Adaptive choice of input states

**2603.26655**

# Thank you!

Contact: `alicezheng@vt.edu`