# A Quantum Speed-Up for Approximating the Top Eigenvector of a Matrix via Improved Tomography

András Gilyén
HUN-REN Alfréd Rényi Institute of Mathematics

Joint work with Yanlin Chen and Ronald de Wolf

Simons Institute, Berkeley, 24th April 2024

# Goal: finding the top eigenvector of a Hermitian $M$

- Used for: computing PageRank, ground states, dimensionality reduction, etc.
- Assume we can query matrix elements of the (dense) matrix $M \in \mathbb{C}^{d\times d}$.
- Full diagonalization takes $\sim d^\omega$ time where $\omega \in [2, 2.37\ldots)$; in practice $\omega \approx 3$.
- A matrix-vector multiplication takes merely $d^2$ time $\Rightarrow$ use power method!

## Power method

- For simplicity assume $\|M\| = 1$, and let $\Delta$ be the spectral gap of $M$
- Sample a random vector $v$
- Iterate $\sim \frac{\log d}{\Delta}$ times:
  compute and update $v \leftarrow Mv$
- The final vector is close to the top eignevector $\psi^{(1)}$ with high probability

- Initial overlap is $|\langle\psi^{(1)}, v\rangle| \approx \frac{1}{\sqrt{d}}$:
  gets magnified by $\sim (1 + \Delta)$ per step.
- The overall time complexity is $\sim d^2/\Delta$; this is provably optimal for $\Delta = \Theta(1)$
- The Lánczos algorithm improves the gap dependence quadratically $\sim d^2/\sqrt{\Delta}$

# Speed up using quantum linear algebra?

- Preparing $|v\rangle$ takes time $\mathcal{O}(\log d)$ assuming QRAM (Kerenidis-Prakash'17)
- Applying $M$ takes time $d^{0.5+o(1)}$ (Low'18 + **G**, Su, Low, Wiebe'18)
- Applying $\Pi = |\psi^{(1)}\rangle\langle\psi^{(1)}|$ takes time $d^{0.5+o(1)}/\Delta$ (QSVT – **G**, Su, Low, Wiebe'18)
- Amplifying $\Pi v$ to $\Pi v/\|\Pi v\|$ has overhead $\sim \sqrt{d}$
- Tomography of $\Pi v$ has overhead $\sim d/\varepsilon$ (Apeldoorn, Cornelissen, **G**, Nannicini'22)
- Combining everything the total complexity is $\sim d^2/(\Delta\varepsilon)$
- We prove a quantum query lower bound $\sim d^{1.5}$ when $\Delta, \varepsilon = \Theta(1)$
- There is some hope, the running time is $\sim d^{1.5}/(|\langle\psi^{(1)}, v\rangle|\Delta\varepsilon)$

# Needle in the haystack: quantum-classical conversions

- Classical amplification is for free — skip expensive quantum amplification
- Do tomography directly on $|0\rangle|\Pi v\rangle + |1\rangle\big|(I - \Pi)v\rangle$
- Power method is robust to small errors $\Pi v + \zeta$ (Hardt and Price'14)
- Suffices to ensure the error term is $\|\zeta\| \le \varepsilon$ and $|\langle\psi^{(1)}, \zeta\rangle| \le \frac{\varepsilon}{\sqrt{d}}$ in each iteration
- We develop a new pure-state tomography procedure with such guarantees
- The complexity of our tomography algorithm remains $\sim d/\varepsilon$ (essentially optimal)

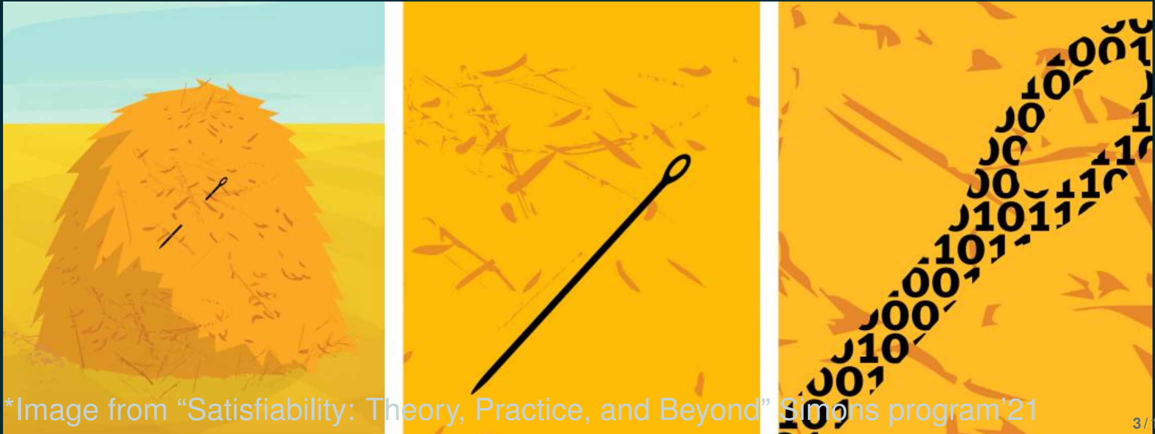

*Image from "Satisfiability: Theory, Practice, and Beyond" Simons program'21

# Tomography by computational basis measurements $|\cdot|$

## Okamoto-Hoeffding Bound

Let $0 \leq X \leq 1$ be a bounded random variable, and $s := \frac{X^{(1)}+X^{(2)}+\ldots+X^{(n)}}{n}$ be the empirical mean of $n$ i.i.d. samples. Then for $p := \mathbb{E}[X]$ we have that

$$\mathbb{P}(\sqrt{s} \geq \sqrt{p} + \varepsilon) \leq \exp\left(-2\varepsilon^2 n\right),$$
$$\mathbb{P}(\sqrt{s} \leq \sqrt{p} - \varepsilon) \leq \exp\left(-\varepsilon^2 n\right).$$

## Gate-efficient estimation of absolute amplitudes

Given $\frac{1}{\varepsilon^2}\ln\left(\frac{2d}{\delta}\right)$ samples of the pure quantum state $|\varphi\rangle := |0\rangle|\psi\rangle + |1\rangle|\cdot\rangle \in \mathbb{C}^{2d}$, measure each copy and let $s_i$ be the frequency of $0, i$ outcomes and define

$$\bar{\psi}_i := \sqrt{s_i}.$$

With probability at least $1-\delta$ it gives an $\varepsilon$-$\ell_\infty$ approximation of the absolute vales $|\psi|$.

# Tomography using conditional samples ($\mathbb{R}$)

**Gate-efficient tomography of "real" states using reference state $|u\rangle = \sum_{i=0}^{d-1} \frac{1}{\sqrt{d}}|i\rangle$**

Given $\frac{4d}{\varepsilon^2}\ln\left(\frac{4d}{\delta}\right)$ samples of the state $|\varphi\rangle := (|+\rangle(|0\rangle|\psi\rangle + |1\rangle|\cdot\rangle) + |-\rangle|0\rangle|u\rangle)/\sqrt{2} \in \mathbb{R}^{4d}$, measure each copy and let $s_{b,i}$ be the frequency of $b, 0, i$ outcomes and define

$$\tilde{\psi}_i := \begin{cases} \max\{0, \quad 2\sqrt{s_{0,i}} - \frac{1}{\sqrt{d}}\} & \text{if } s_{0,i} > s_{1,i} \\ 0 & \text{if } s_{0,i} = s_{1,i} \\ \min\{0, -(2\sqrt{s_{1,i}} - \frac{1}{\sqrt{d}})\} & \text{if } s_{0,i} < s_{1,i} \end{cases}$$

With probability at least $1-\delta$ it gives an $\frac{\varepsilon}{\sqrt{d}}$-$\ell_\infty$ (and thus $\varepsilon$-$\ell_2$) approximation of $\psi$.

- Idea can be extended for $\psi \in \mathbb{C}^d$, see (Apeldoorn, Cornelissen, **G**, Nannicini'22).

# Unbiased tomography using a reference state $\bar{\psi} \in \mathbb{R}^d$

Suppose we have a reference state $\bar{\psi} \in \mathbb{R}^d$ such that $|\bar{\psi}_j|^2 \geq \max\{\frac{\varepsilon^2}{d}, \frac{2}{3}|\psi_j|^2\}$ for all $j \in [d]$. Given $\frac{12d}{\varepsilon^2} \ln\left(\frac{8d}{\delta}\right)$ copies of the pure state $|\varphi\rangle := \left(|+\rangle|\psi\rangle + |-\rangle\left|\bar{\psi}\right\rangle\right)/\sqrt{2} \in \mathbb{C}^d$, measure each copy and let $s_{b,i}$ be the frequency of $b, i$ outcomes, then

$$\tilde{\psi}_j := \frac{s_{0,j} - s_{1,j}}{\bar{\psi}_j}$$

is an unbiased estimator of $\Re(\psi)$. Moreover, for any $B = \{v^{(j)} : j \in [k]\}$ ONB we have

$$\Pr\left[\forall v \in B : |\langle\tilde{\psi} - \Re(\psi)|v\rangle| < \frac{\varepsilon}{\sqrt{d}}\right] \geq 1 - \delta.$$

$$|\langle 0, i|\phi\rangle|^2 - |\langle 1, i|\phi\rangle|^2 = \left|\frac{\psi_i + \bar{\psi}_i}{2}\right|^2 - \left|\frac{\psi_i - \bar{\psi}_i}{2}\right|^2 = \Re(\psi_i)\bar{\psi}_i$$

Concentration follows from the Bennett-Bernstein Bound as $\frac{1}{|\bar{\psi}_i|} \leq \frac{\sqrt{d}}{\varepsilon}$ and $\|\mathrm{Cov}\| \leq 1$.

# Application: Tomography using a reflection $2|\psi\rangle\langle\psi| - I$

## Algorithm – quantum noisy power method – Chen, G, de Wolf [QIP'24]

- Input: Controlled reflection $R = 2|\psi\rangle\langle\psi| - I$ or block-encoding of $\Pi = |\psi\rangle\langle\psi|$
- Init: sample Gaussian random vector $\phi^{(0)}$
- For $j = 0 \ldots \log(d)$ do

  Prepare data structure in QRAM for preparing the normalized state $\left|\phi^{(j)}\right\rangle$

  Do tomography on $\Pi\left|\phi^{(j)}\right\rangle$ to $\ell_2$-precision $\varepsilon$ giving $\phi^{(j+1)}$

- Initial overlap with $|\psi\rangle$ is $\sim \frac{1}{\sqrt{d}}$, error overlap is $\sim \frac{\varepsilon}{\sqrt{d}}$ in each iteration
- In each iteration the overlap doubles until it is $\Omega(1)$ – total complexity is $\sim \frac{d\log(d)}{\varepsilon}$

- Given $U$ block-encoding of a matrix $M$ having gap $\Delta$ turn into top-eigenvector projector using QSVT by $\sim \frac{1}{\Delta}$ iterations
- Sparse $M$ can be block-encoded by $\sim \sqrt{s}$ queries $\Rightarrow \sim d\sqrt{s}/(\Delta\varepsilon)$ overall
- Query complexity is $\sim d^{1.5}/\Delta$ for dense case and $\varepsilon = \Theta(1)$.

# Extension: Process tomography of reflections $2\Pi - I$

- Similar algorithm works if we are promised rank $\Pi$ is at most $r$
- Sample $\sim r$ independent random vectors and apply power method for each
- The subspace spanned by the left singular vectors with s.v. $\Omega(1)$ approximate $\Pi$
- Algorithm uses $R = 2\Pi - I$ about $\sim \frac{dr}{\varepsilon}$ times
- Probably optimal, we prove $\widetilde{\Omega}(dr + \frac{d}{\varepsilon})$ lower bound
- Unitary tomography has complexity $\Theta(d^2/\varepsilon)$ (Haah, Kothari, O'Donnell, Tang'23)

## Application to sparse matrices – Chen, G, de Wolf [QIP'24]

- Given $U$ block-encoding of a matrix $M$ having gap $\gamma$ below top-$r$ eigensubspace turn into top-eigenvector projector using QSVT by $\sim \frac{1}{\Delta}$ iterations
- Sparse $M$ can be block-encoded by $\sim \sqrt{s}$ queries $\Rightarrow \mathcal{O}\big(dr\sqrt{s}/(\Delta\varepsilon)\big)$ overall

# Iterative refinement

- Comes from early days of classical computing having limited precision numbers
- Solves a large linear equation system given such limited arithmetic precision
- Idea is to solve it only to constant precision, and then recurse
  - Compute $\tilde{x}$ such that $\left\|A\tilde{x} - b\right\| \leq \frac{1}{2}\|b\|$
  - Set $b \leftarrow A\tilde{x} - b$ and repeat $\log_2(1/\varepsilon)$ times
  - Take the sum $\bar{x} = \sum \tilde{x}$ of the intermediate solutions
  - The result satisfies $\|A\bar{x} - b\| \leq \varepsilon\|b\|$
- In the quantum case this requires updating the state preparation unitary
- E.g., use classical write quantum read QRAM
- Could get a polynomial speedup when classical output is required (IP solvers)
- Idea pioneered by Mohammadhossein Mohammadisiahroudi, Brandon Augustino, Tamás Terlaky, et al.
- Similar ideas can be applied to pure state tomography

# Iteratively refined tomography

## Refinement step

- Want to learn $\psi^{(0)}$, by improving the current estimate $\psi^{(1)}$
- Input: $\varepsilon \in (0,2]$, unitaries $U^{(0)}, U^{(1)}$ such that $\left\|\psi^{(0)} - \psi^{(1)}\right\| \leq \varepsilon$ for $\psi^{(i)} := (\langle 0^a| \otimes I)U^{(i)}|0^q\rangle$
- Let $W := |+\rangle\langle+| \otimes U^{(0)} - |-\rangle\langle-| \otimes U^{(1)}$, and $k \approx 1/\varepsilon$
- Let $AA(W,k)$ be the $k$-step amplitude amplification of $|\varphi\rangle := \psi^{(0)} - \psi^{(1)} = (\langle 0^{a+1}| \otimes I)W|0^{q+1}\rangle$
- Perform tomography on $|\phi\rangle := (\langle 0^{a+1}| \otimes I)AA(W,k)|0^{q+1}\rangle$ to $\ell_2$-precision $\frac{1}{6}$ giving estimate $\tilde{\phi}$
- Output: $\psi^{(1)} \leftarrow \psi^{(1)} + \frac{2\tilde{\phi}}{2k+1}$ (satisfying $\left\|\psi^{(0)} - \psi^{(1)}\right\| \leq \varepsilon/2$ with probability $\geq 1-\delta$)
- Complexity: $\approx \frac{1}{\varepsilon} \times d$

# Quantum lower bound

## Hard instance

- Hide $d$ bits of information in $\phi \in \{-1,+1\}^d$.
- Define $M := \frac{|\phi\rangle\langle\phi|}{2d} + \mathcal{N}(0, \frac{1}{100\sqrt{d}})$ entry-wise noise.
- Because of random matrix theory with high probability $\|M\| \leq 1$ and $\Delta \geq \frac{1}{4}$.
- The top-eigenvector $\psi^{(1)}$ is $\mathcal{O}(1)$-close to $\phi$. Learning $\psi^{(1)}$ reveals $\Omega(d)$ bits.

## Quantum lower bound – using a variant of the adversary method

- To learn $\phi_i$ we need to extract sign $\pm\frac{1}{d}$ from noisy entries $\pm\frac{1}{d} + \mathcal{N}(0, \frac{1}{100\sqrt{d}})$.
- Requires $\sim d$ "samples" classically $\Rightarrow \sim \Omega(d^2)$ queries to learn most bits.
- Requires $\sim \sqrt{d}$ "queries" quantumly $\Rightarrow \sim \Omega(d^{1.5})$ queries to learn most bits.