Near Optimality of the Lanczos Method for Matrix Functions

Christopher Musco, New York University

Joint work with:

Noah Amsel, New York University
Tyler Chen, New York University
Anne Greenbaum, University of Washington
Cameron Musco, UMass Amherst

THE LANCZOS METHOD

The Lanczos Method is a single algorithm that underlies state of the art iterative methods for:

- solving linear systems,
- approximating eigenvectors and eigenvalues,
- approximating matrix functions,
- and much more.

Introduced in 1950, developed through the 70s, ubiquitous in well-developed scientific computing libraries.

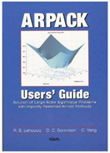

Meta-observation: Despite decades of very good theoretical work, for a wide range of problems, the Lanczos method often performs far better than our best theory predicts.

1. Converges faster than expected.
2. Is more robust to round-off error on finite precision computers than expected.

Meta-observation: Despite decades of very good theoretical work, for a wide range of problems, the Lanczos method often performs <u>far better</u> than our best theory predicts.

1. Converges faster than expected.
2. Is more robust to round-off error on finite precision computers than expected.

[Musco, Musco, Sidford, SODA 2018]
[Chen, Greenbaum, Musco, Musco, SIMAX 2022]
[Amsel, Chen, Musco, Musco, Greenbaum, 2023]
[Meyer, Musco, Musco, SODA 2024]
[Bhattacharjee, Jayaram, Musco, Musco, Ray, 2024]

Meta-observation: Despite decades of very good theoretical work, for a wide range of problems, the Lanczos method often performs far better than our best theory predicts.

1. Converges faster than expected.
2. Is more robust to round-off error on finite precision computers than expected.

[Musco, Musco, Sidford, SODA 2018]
[Chen, Greenbaum, Musco, Musco, SIMAX 2022]
[Amsel, Chen, Musco, Musco, Greenbaum, 2023]
[Meyer, Musco, Musco, SODA 2024]
[Bhattacharjee, Jayaram, Musco, Musco, Ray, 2024]

Today: focus on Lanczos for matrix function approximation.

WHAT IS A MATRIX FUNCTION?

For today, just consider symmetric matrices $\mathbf{A} \in \mathbb{R}^{d \times d}$, which always have an eigendecomposition:

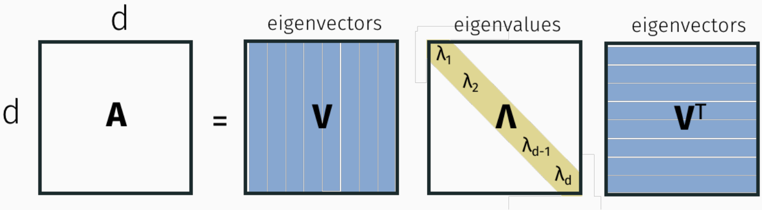

where $\mathbf{V}$ is orthogonal and $\lambda_1, \ldots, \lambda_n$ are real.

For any scalar function $f : \mathbb{R} \to \mathbb{R}$ define $f(\mathbf{A})$:

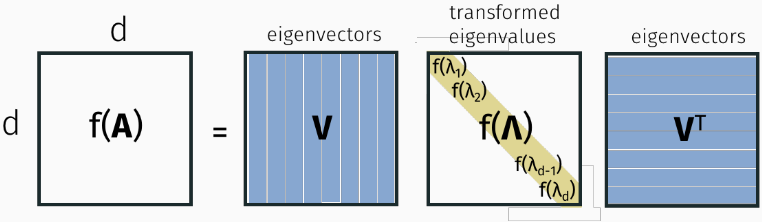

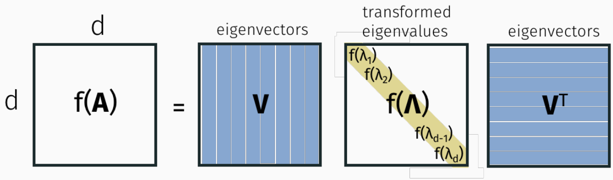

- When $f(x) = \frac{1}{x}$, $f(\mathbf{A}) = \mathbf{A}^{-1}$. $f(\mathbf{A})\mathbf{b}$ solves the system, $\mathbf{A}\mathbf{x} = \mathbf{b}$.
- When $\mathbf{A}$ has non-negative eigenvalues and $f(x) = \sqrt{x}$, $f(\mathbf{A})$ is the matrix square root. $f(\mathbf{A})\mathbf{g}$ samples a multivariate Gaussian vector with covariance $\mathbf{A}$.
- The matrix exponential, $f(x) = e^x$, finds applications in differential equations, control theory, computational chemistry, combinatorial optimization, and more.

APPLICATIONS OF MATRIX FUNCTIONS

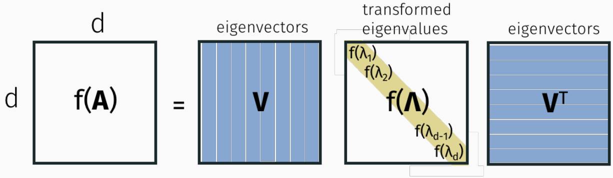

Other important matrix functions: log, absolute value, sign function, window functions, inverse square root, etc.

In many cases, $\text{tr}(f(\mathbf{A}))$ is a meaningful quantity. E.g., $\text{tr}(\mathbf{A}^q)$ can be used to count cycles in a graph adjacency matrix. $\text{tr}(\log(\mathbf{A}))$ is the log determinant. The trace of a window function applied to $\mathbf{A}$ counts the number of eigenvalues in a given interval.

Cost to compute $f(\mathbf{A})$:

$$\underbrace{O(n^3)}_{\text{eigendecompose } \mathbf{A} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T} + \underbrace{O(n)}_{\text{compute } f(\boldsymbol{\Lambda})} + \underbrace{O(n^3)}_{\text{form } \mathbf{V} f(\boldsymbol{\Lambda})\, \mathbf{V}^T}$$

$$= O(n^3).$$

Cost to compute $f(\mathbf{A})$:

$$\underbrace{O(n^3)}_{\text{eigendecompose } \mathbf{A} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T} + \underbrace{O(n)}_{\text{compute } f(\boldsymbol{\Lambda})} + \underbrace{O(n^3)}_{\text{form } \mathbf{V} f(\boldsymbol{\Lambda})\mathbf{V}^T}$$

$$= O(n^3).$$

In theory, can be improved to $O(n^\omega) \approx O(n^{2.371866})$.
(but this is still slow)

Typically only interested in computing $f(\mathbf{A})\mathbf{b}$ for some $\mathbf{b} \in \mathbb{R}^n$.

Even for $\text{tr}(f(\mathbf{A}))$, this is true, since we can estimate trace via the identity $\text{tr}(f(\mathbf{A})) = \mathbb{E}[\mathbf{g}^T f(\mathbf{A})\mathbf{g}]$ (Hutchinson's estimator).

$$f\left(\begin{bmatrix} & & \\ & \mathbf{A} & \\ & & \end{bmatrix}\right) \cdot \begin{bmatrix} \\ \mathbf{b} \\ \\ \end{bmatrix}$$

Typically only interested in computing $f(\mathbf{A})\mathbf{b}$ for some $\mathbf{b} \in \mathbb{R}^n$.

Even for $\text{tr}(f(\mathbf{A}))$, this is true, since we can estimate trace via the identity $\text{tr}(f(\mathbf{A})) = \mathbb{E}[\mathbf{g}^T f(\mathbf{A})\mathbf{g}]$ (Hutchinson's estimator).

$$f\left(\begin{bmatrix} & & \\ & \mathbf{A} & \\ & & \end{bmatrix}\right) \cdot \begin{bmatrix} \\ \mathbf{b} \\ \\ \end{bmatrix}$$

Often much cheaper than computing $f(\mathbf{A})$ explicitly!

Krylov subspace methods are the dominant approach for approximating $f(\mathbf{A})\mathbf{b}$ in less than $O(n^3)$ time.

Key observation: Low degree matrix **polynomials** can be computed efficiently.

$$p\left(\begin{bmatrix} & & \\ & A & \\ & & \end{bmatrix}\right) \cdot \begin{bmatrix} \\ b \\ \\ \end{bmatrix}$$

KRYLOV SUBSPACE METHODS

Key observation: Low degree matrix **polynomials** can be computed efficiently.

$$p\left(\begin{bmatrix} & A & \end{bmatrix}\right) \cdot \begin{bmatrix} b \end{bmatrix}$$

$$V\Lambda^k V^T b = V\Lambda V^T V\Lambda V^T \cdots V\Lambda V^T b = A^k b$$

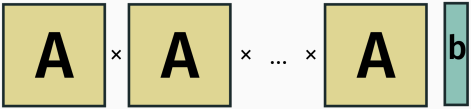

Key observation: Low degree matrix **polynomials** can be computed efficiently.

$$\mathbf{p}\left(\begin{bmatrix} & \mathbf{A} & \end{bmatrix}\right) \cdot \begin{bmatrix} \mathbf{b} \end{bmatrix}$$

$$\mathbf{V}\boldsymbol{\Lambda}^k\mathbf{V}^T\mathbf{b} = \mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T\cdots\mathbf{V}\boldsymbol{\Lambda}\mathbf{V}^T\mathbf{b} = \mathbf{A}^k\mathbf{b}$$

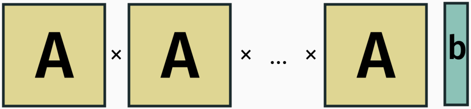

Total time to compute $p(\mathbf{A})\mathbf{b} = c_0\mathbf{b} + c_1\mathbf{A}\mathbf{b} + c_2\mathbf{A}^2\mathbf{b} + \ldots + c_k\mathbf{A}^k\mathbf{b}$:

$$O(k \cdot n^2) \ll O(n^3).$$

For general matrix functions: approximate $f(x)$ with low-degree polynomial $p(x)$ so $f(\mathbf{A})\mathbf{b} \approx p(\mathbf{A})\mathbf{b}$.

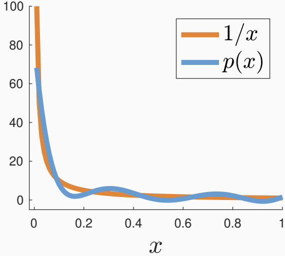

The **Lanczos method** gives one particular way of doing this that works for <u>any</u> function f. When $\mathbf{A}$ is positive definite, and $f(x) = 1/x$, it is equivalent to the **Conjugate Gradient** method.

Other Krylov subspace methods: MINRES, Richardson iteration / gradient descent, accelerated gradient descent, etc.

$$\|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \leq \|f(\mathbf{A}) - p(\mathbf{A})\| \cdot \|\mathbf{b}\| \leq \epsilon \cdot \|\mathbf{b}\|$$

where

$$\epsilon = \max_{i=1,\ldots,n} |f(\lambda_i) - p(\lambda_i)|.$$

$$\|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \leq \|f(\mathbf{A}) - p(\mathbf{A})\| \cdot \|\mathbf{b}\| \leq \epsilon \cdot \|\mathbf{b}\|$$

where

$$\epsilon = \max_{i=1,\ldots,n} |f(\lambda_i) - p(\lambda_i)|.$$

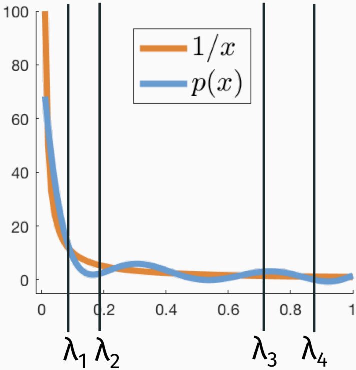

$$\|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \leq \|f(\mathbf{A}) - p(\mathbf{A})\| \cdot \|\mathbf{b}\| \leq \epsilon \cdot \|\mathbf{b}\|$$

where

$$\epsilon = \max_{i=1,\ldots,n} |f(\lambda_i) - p(\lambda_i)|.$$

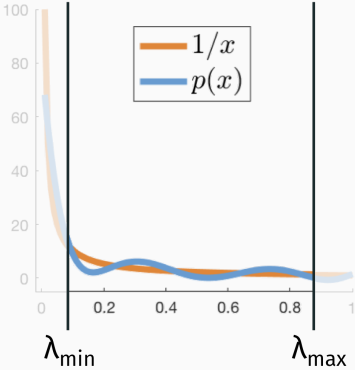

If we know $\lambda_{\min}(\mathbf{A})$ and $\lambda_{\max}(\mathbf{A})$ we can explicitly compute an optimal polynomial p for uniformly approximating f.

$$\delta_k = \min_{\text{degree } k \text{ poly } p} \left(\max_{x \in [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]} |f(x) - p(x)| \right)$$

Final bound: Return $p(\mathbf{A})\mathbf{b}$ such that

$$\|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \leq \delta_k \cdot \|\mathbf{b}\|.$$

Example bounds:

- Linear systems in $k = O\left(\sqrt{\lambda_{\max} / \lambda_{\min}} \log(1/\epsilon)\right)$ iterations.
- Matrix sign function in $k = O(1/\epsilon)$ iterations.
- Top eigenvector in $k = O(\log(n)/\sqrt{\epsilon})$ iterations.

Example bounds:

- Linear systems in $k = O\left(\sqrt{\lambda_{\max} / \lambda_{\min}} \log(1/\epsilon)\right)$ iterations.
- Matrix sign function in $k = O(1/\epsilon)$ iterations.
- Top eigenvector in $k = O(\log(n)/\sqrt{\epsilon})$ iterations.

But, we need to know λ_{min} and λ_{max}, and finding/representing an optimal p can be challenging.

The Lanczos method avoids these issues and performs much better in practice.

Step 1: Form orthogonal matrix $\mathbf{Q} = [\mathbf{q}_0, \mathbf{q}_1, \ldots, \mathbf{q}_k]$ that spans the Krylov subspace

$$\mathcal{K} = \{\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \ldots \mathbf{A}^k\mathbf{b}\}.$$

Step 1: Form orthogonal matrix $\mathbf{Q} = [\mathbf{q}_0, \mathbf{q}_1, \ldots, \mathbf{q}_k]$ that spans the Krylov subspace

$$\mathcal{K} = \{\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \ldots \mathbf{A}^k\mathbf{b}\}.$$

Step 2: Compute (tridiagonal)

$$\mathbf{T} = \mathbf{Q}^T\mathbf{A}\mathbf{Q}$$

Step 1: Form orthogonal matrix $\mathbf{Q} = [\mathbf{q}_0, \mathbf{q}_1, \dots, \mathbf{q}_k]$ that spans the Krylov subspace

$$\mathcal{K} = \{\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}, \dots \mathbf{A}^k\mathbf{b}\}.$$

Step 2: Compute (tridiagonal)

$$\mathbf{T} = \mathbf{Q}^T\mathbf{A}\mathbf{Q}$$

Step 3: Approximate $f(\mathbf{A})\mathbf{b}$ by

$$\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}$$

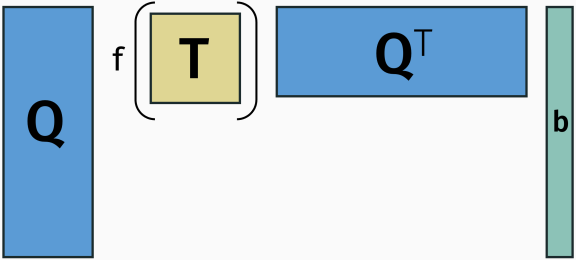

Runtime: $O(n^2k + nk + k^2 \log k)$

Reduce the problem to the cost of computing a matrix function for a $k \times k$ matrix.

Current state-of-the-art convergence result for Lanczos:

Theorem (Implicit in Saad, '92)

Let $\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T$ *be the output of Lanczos run on* $\mathbf{A}, \mathbf{b}$ *for* k *iterations with function* f*. Then, for any* f*:*

$$\|\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b} - f(\mathbf{A})\mathbf{b}\| \leq 2 \cdot \delta_k \cdot \|\mathbf{b}\|,$$

where

$$\delta_k = \min_{\text{degree } k \text{ poly } p} \left(\max_{x \in [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]} |f(x) - p(x)| \right).$$

Takeaway: Lanczos matches the best uniform polynomial approximation up to a factor of two! And we didn't even need to do any computation involving polynomials.

Very powerful result with straightforward proof!

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

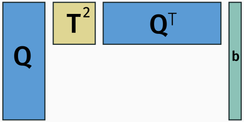

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

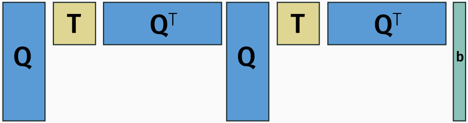

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

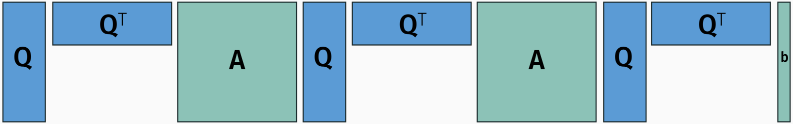

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

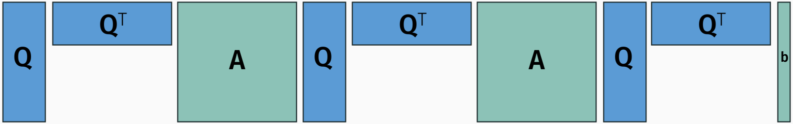

$\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}$ all lie in the span of $\mathbf{Q}$ (the degree k Krylov subspace).

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

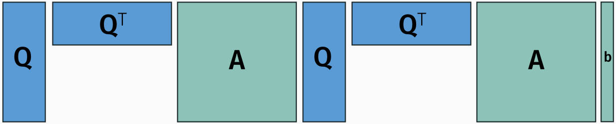

$\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}$ all lie in the span of $\mathbf{Q}$ (the degree k Krylov subspace).

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

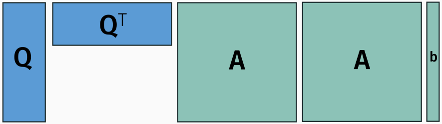

$\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}$ all lie in the span of $\mathbf{Q}$ (the degree k Krylov subspace).

Very powerful result with straightforward proof!

Claim 1: Lanczos applies degree k polynomials exactly.

Proof:

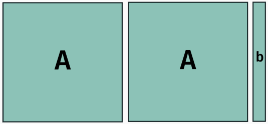

$\mathbf{b}, \mathbf{A}\mathbf{b}, \mathbf{A}^2\mathbf{b}$ all lie in the span of $\mathbf{Q}$ (the degree k Krylov subspace).

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial “part” of f.

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \leq{}& \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\ &+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\|\end{aligned}$$

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| &\leq \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\ &+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &\leq \delta_k\|\mathbf{b}\|\end{aligned}$$

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}
\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| &\leq \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\
&+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\
&\leq \delta_k\|\mathbf{b}\| + 0
\end{aligned}$$

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| &\leq \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\ &+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &\leq \delta_k\|\mathbf{b}\| + 0 + \|p(\mathbf{T}) - f(\mathbf{T})\| \cdot \|\mathbf{Q}^T\mathbf{b}\|\end{aligned}$$

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}
\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| &\leq \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\
&+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\
&\leq \delta_k\|\mathbf{b}\| + 0 + \|p(\mathbf{T}) - f(\mathbf{T})\| \cdot \|\mathbf{Q}^T\mathbf{b}\|
\end{aligned}$$

Since $\mathbf{T} = \mathbf{Q}^T\mathbf{A}\mathbf{Q}$, $[\lambda_{\min}(\mathbf{T}), \lambda_{\max}(\mathbf{T})] \subseteq [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$.

For general functions $f(x)$ we use triangle inequality to show that Lanczos automatically applies the polynomial "part" of f.

Let p be the optimal degree k polynomial approximation to f on $[\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$:

$$\begin{aligned}\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| &\leq \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\ &+ \|p(\mathbf{A})\mathbf{b} - \mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &+ \|\mathbf{Q}p(\mathbf{T})\mathbf{Q}^T\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \\ &\leq \delta_k\|\mathbf{b}\| + 0 + \delta_k\|\mathbf{b}\|.\end{aligned}$$

Since $\mathbf{T} = \mathbf{Q}^T\mathbf{A}\mathbf{Q}$, $[\lambda_{\min}(\mathbf{T}), \lambda_{\max}(\mathbf{T})] \subseteq [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$.

Theorem (Implicit in Saad, '92)

Let $\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T$ be the output of Lanczos run on $\mathbf{A}, \mathbf{b}$ for k iterations with function f. Then, for any f:

$$\|\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b} - f(\mathbf{A})\mathbf{b}\| \leq 2 \cdot \delta_k \cdot \|\mathbf{b}\|,$$

where

$$\delta_k = \min_{\text{degree } k \text{ poly } p} \left(\max_{x \in [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]} |f(x) - p(x)| \right).$$

Really great bound, but is this the end of the story?

Lanczos almost always performs even better than the uniform convergence bound predicts. Often by orders of magnitude.

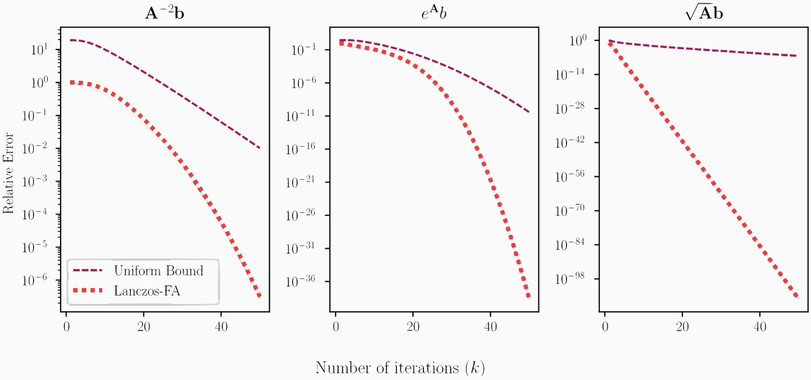

What is the right bound?

Conjecture (Instance Optimality of Lanczos)

For a wide-variety of matrix functions, the Lanczos method performs nearly as well as the best solution in the Krylov subspace. I.e., for some approximation factor C,

$$\|\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T - f(\mathbf{A})\mathbf{b}\| \leq C \cdot \min_{\textit{degree } k \textit{ poly } p} \|f(\mathbf{A}) - p(\mathbf{A})\|.$$

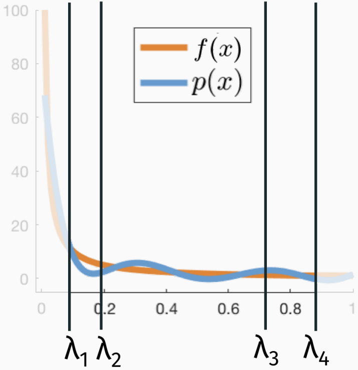

I.e., we believe Lanczos is competitive with polynomials that are only accurate at A's eigenvalues, instead of on the entires interval $[\lambda_{\text{min}}(\mathbf{A}), \lambda_{\text{max}}(\mathbf{A})]$. Despite the fact that it doesn't have enough information to compute A's eigenvalues.

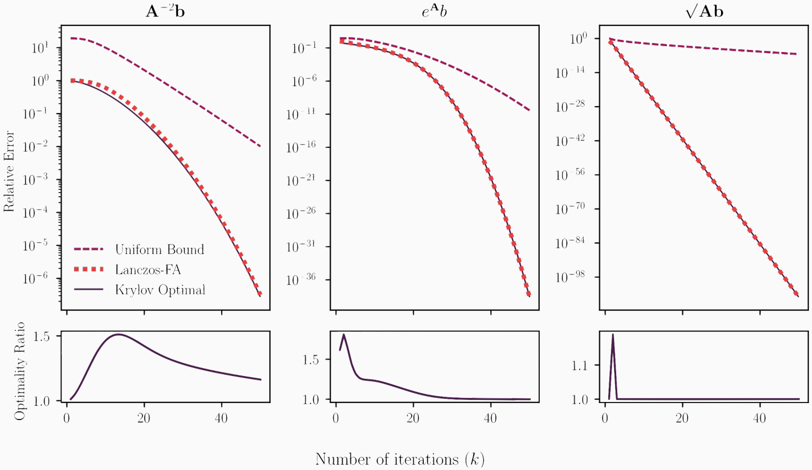
$\mathbf{A}^{-2}\mathbf{b}$
$e^{\mathbf{A}}b$
$\sqrt{\mathbf{A}}\mathbf{b}$
Relative Error
Optimality Ratio
Uniform Bound
Lanczos-FA
Krylov Optimal
Number of iterations (k)

Conjecture is known to hold for the special case of $f(x) = 1/x$ when A is positive definite.

Claim (Optimality of Lanczos/CG for Linear Systems)

For any positive definite $\mathbf{A}$,

$$\|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T - f(\mathbf{A})\mathbf{b}\|_{\mathbf{A}} = \min_{\text{degree } k \text{ poly } p} \|f(\mathbf{A}) - p(\mathbf{A})\|_{\mathbf{A}}$$

As a consequence, letting $\kappa(\mathbf{A}) = \lambda_{\min}(\mathbf{A}) / \lambda_{\max}(\mathbf{A})$ *be the condition number of* $\mathbf{A}$,

$$\|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T - f(\mathbf{A})\mathbf{b}\| \leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|f(\mathbf{A}) - p(\mathbf{A})\|.$$

A related but weaker guarantee was shown for the matrix exponential by [Druskin, Greenbaum, Knizhnerman ‘98], but otherwise no near-optimality guarantees are known for any other functions.

Lanczos is near-optimal for rational functions more broadly!

Lanczos is near-optimal for rational functions more broadly!

Setting:

- Let $r(x) = \frac{(x-w_1)(x-w_2)\ldots(x-w_m)}{(x-z_1)(x-z_2)\ldots(x-z_q)}$ be a degree-(m,q) rational function with real poles lying outside $\mathbf{A}$'s spectral range. I.e., $z_1, \ldots, z_q \notin [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$.

Lanczos is near-optimal for rational functions more broadly!

Setting:

- Let $r(x) = \frac{(x-w_1)(x-w_2)\ldots(x-w_m)}{(x-z_1)(x-z_2)\ldots(x-z_q)}$ be a degree-(m, q) rational function with real poles lying outside $\mathbf{A}$'s spectral range. I.e., $z_1, \ldots, z_q \notin [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]$.

Theorem (Main result)

Lanczos is near-instance optimal for a such a rational function with $C = q \cdot \prod_{i=1}^{q} \kappa(\mathbf{A} - z_i \mathbf{I})$. Specifically, for $k \geq \max\{m, q-1\}$,

$$\|f(\mathbf{A})\mathbf{b} - \mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b}\| \leq C \cdot \min_{\text{degree } (k-q+1) \text{ poly. } p} \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.$$

- Our approximation factor $C = q \cdot \prod_{i=1}^{q} \kappa(\mathsf{A} - z_i \mathbf{I})$ is really bad. Grows exponentially in q. We believe it can be significantly improved.
- The worst case empirical value we observed for C when all poles are at 0 is roughly $\sqrt{q \cdot \kappa(\mathsf{A})}$.
- The requirement that $z_1, \ldots, z_q \notin [\lambda_{\min}(\mathsf{A}), \lambda_{\max}(\mathsf{A})]$ is necessary for a true near-optimality bound, but we might hope to prove slightly weaker results. More on this later.

Despite the seeming looseness in our bound, it often more accurately reflects the performance of Lanczos in practice than the classic uniform approximation bound does.

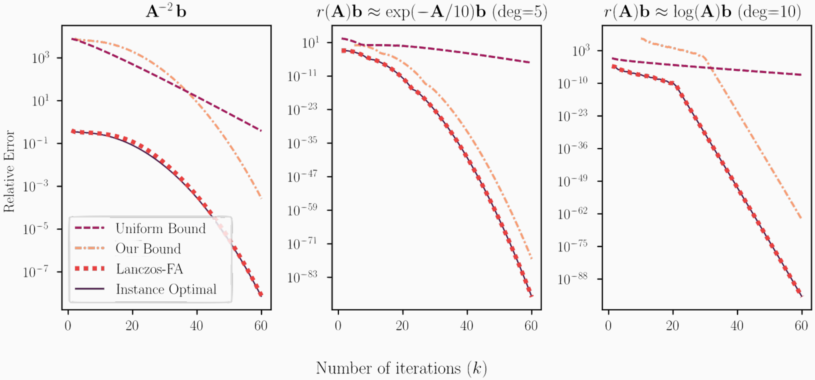

WHY DO WE CARE ABOUT RATIONAL FUNCTIONS?

- Rational functions are interesting in their own right. They include e.g. $1/x$, $1/x^q$, etc.
- More importantly, rational functions often give very accurate approximations to other functions, so **their behavior can tell use about other functions.**
- For example, a uniform polynomial approximation to $\sqrt{x}$ on $[\lambda_{\text{min}}, \lambda_{\text{max}}]$ requires $O(\sqrt{\lambda_{\text{max}} / \lambda_{\text{min}}})$ degree. A uniform rational approximation requires just $O(\log(\lambda_{\text{max}} / \lambda_{\text{min}}))$ degree. Similar improvements are possible to for x^α for other choices of α, $\exp(-x)$, etc.

Convergence for $f(\mathsf{A}) = \mathsf{A}^{-0.4}$.

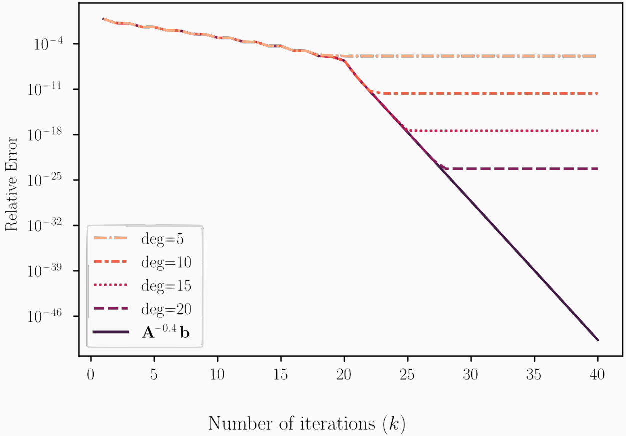

Behavior of Lanczos for $f(\mathsf{A})$ closely tracks behavior for rational approximations of f.

Formally, if we have a C-factor near-optimality result for rational functions of degree (m, q), a simple application of triangle inequality shows that:

$$\|\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b} - f(\mathbf{A})\mathbf{b}\| \leq C \cdot \min_{\text{degree } k \text{ poly } p} \|f(\mathbf{A})\mathbf{b} - p(\mathbf{A})\mathbf{b}\| + (C+2) \cdot \gamma_{m,q} \cdot \|\mathbf{b}\|_2,$$

where $\gamma_{m,q}$ is error of the optimal degree-(m, q) rational approximation to f on $[\lambda_{\min}, \lambda_{\max}]$.

So, Lanczos is near optimal for f, up to a term depending on the error of the best uniform rational approximation. Typically far smaller than the error of the best uniform polynomial approximation that appears in current bounds for Lanczos.

Our proof starts with the instance optimality of Lanczos (equivalently CG) for applying $f(x) = 1/x$. I.e., $f(\mathbf{A})\mathbf{b} = \mathbf{A}^{-1}\mathbf{b}$.

$$\|\mathbf{A}^{-1}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{b}\| \leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-1}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|$$

Follows from the fact that Lanczos computes the $\mathbf{A}$-norm optimal approximation to $\mathbf{A}^{-1}\mathbf{b}$ in the Krylov subspace.

In particular, the Krylov subspace is spanned by $\mathbf{Q}$. To project a vector $\mathbf{y}$ onto $\mathbf{Q}$ in the $\mathbf{A}$-norm, $\|\cdot\|_{\mathbf{A}}$, we apply the projector:

$$\mathbf{Q}(\mathbf{Q}^T\mathbf{A}\mathbf{Q})^{-1}\mathbf{Q}^T\mathbf{A}\mathbf{y}.$$

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathsf{A}) = \mathsf{A}^{-2}$.

PROOF SKETCH

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

PROOF SKETCH

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 1: $\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$.

- By previous slide, $\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}$ is the best approximation to $\mathbf{A}^{-2}\mathbf{b}$ in the span of the Krylov subspace in the $\mathbf{A}$-norm. So we have:

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{p}(\mathbf{A})\mathbf{b}\|.$$

PROOF SKETCH

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 2: $\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 2: $\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$

$$\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \leq \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\| \cdot \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{b} - \mathbf{A}^{-1}\mathbf{b}\|$$

PROOF SKETCH

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 2: $\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$

$$\begin{aligned}\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| &\leq \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\| \cdot \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{b} - \mathbf{A}^{-1}\mathbf{b}\| \\ &\leq \frac{\sqrt{\kappa(\mathbf{A})}}{\lambda_{\min}(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-1}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.\end{aligned}$$

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 2: $\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$

$$\begin{aligned}\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| &\leq \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\| \cdot \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{b} - \mathbf{A}^{-1}\mathbf{b}\| \\ &\leq \frac{\sqrt{\kappa(\mathbf{A})}}{\lambda_{\min}(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-1}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.\end{aligned}$$

Key Idea: The optimal error for approximating $\mathbf{A}^{-1}$ with degree k can be bounded by the optimal error for approximating $\mathbf{A}^{-2}$ with degree $k-1$. Since $\|\mathbf{A}^{-1}\mathbf{b} - \mathbf{A}p(\mathbf{A})\mathbf{b}\| \leq \lambda_{\max}(\mathbf{A}) \cdot \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|$.

To get a sense of how to generalize this to rational functions, let's consider the special case of $r(x) = 1/x^2$. I.e., $r(\mathbf{A}) = \mathbf{A}^{-2}$.

$$\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| = \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|.$$

Term 2: $\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\|$

$$\begin{aligned}\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| &\leq \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\| \cdot \|\mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{b} - \mathbf{A}^{-1}\mathbf{b}\| \\ &\leq \frac{\sqrt{\kappa(\mathbf{A})}}{\lambda_{\min}(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-1}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.\end{aligned}$$

Key Idea: The optimal error for approximating $\mathbf{A}^{-1}$ with degree k can be bounded by the optimal error for approximating $\mathbf{A}^{-2}$ with degree $k-1$. Since $\|\mathbf{A}^{-1}\mathbf{b} - \mathbf{A}p(\mathbf{A})\mathbf{b}\| \leq \lambda_{\max}(\mathbf{A}) \cdot \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|$.

Overall, this gives:

$$\|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \leq \kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.$$

PROOF SKETCH

Putting it together, we have:

$$
\begin{aligned}
\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| &= \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \\
&\leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\qquad + \kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\leq 2\kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.
\end{aligned}
$$

Putting it together, we have:

$$\begin{aligned}
\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| &= \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \\
&\leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\quad + \kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\leq 2\kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.
\end{aligned}$$

- This gives our main result in the special case of $r(x) = 1/x^2$.

PROOF SKETCH

Putting it together, we have:

$$\begin{aligned}
\|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b}\| &= \|\mathbf{A}^{-2}\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| + \|\mathbf{Q}\mathbf{T}^{-2}\mathbf{Q}^T\mathbf{b} - \mathbf{Q}\mathbf{T}^{-1}\mathbf{Q}^T\mathbf{A}^{-1}\mathbf{b}\| \\
&\leq \sqrt{\kappa(\mathbf{A})} \cdot \min_{\text{degree } k \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\qquad + \kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\| \\
&\leq 2\kappa(\mathbf{A})^{3/2} \cdot \min_{\text{degree } k-1 \text{ poly } p} \|\mathbf{A}^{-2}\mathbf{b} - p(\mathbf{A})\mathbf{b}\|.
\end{aligned}$$

- This gives our main result in the special case of $r(x) = 1/x^2$.
- The general result follows by iterating these types of ideas to bound the error on higher degree rational functions.

OPEN QUESTIONS

- Tighten our bounds. Our worst numerical example for A^{-q} has $C = \sqrt{q\kappa}$. Our best theoretical upper bound is $C = q\kappa^q$.

- Tighten our bounds. Our worst numerical example for $\mathbf{A}^{-q}$ has $C = \sqrt{q\kappa}$. Our best theoretical upper bound is $C = q\kappa^q$.
- Extend our results to the case when $r(x)$ has poles in $\mathbf{A}$'s spectral range. In this case, Lanczos seems to be oscillate between very bad and near optimal solutions.
- We can explain this when $\mathbf{A}$ is not PSD and $r(x) = 1/x$ by relating the convergence of CG to that of MINRES. Lack a general result.

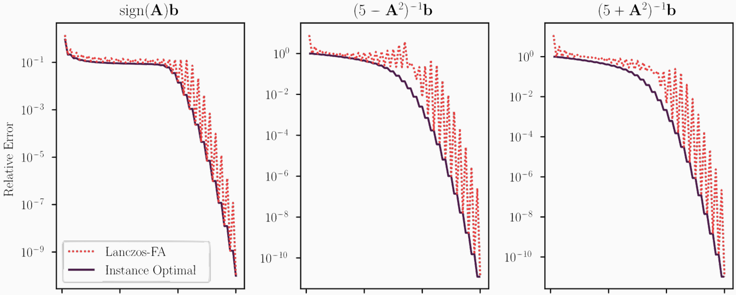

- Prove a direct instance optimality bound for the matrix exponential. Some progress in [Druskin, Greenbaum, Knizhnerman '98].
- Prove instance optimality bounds for the matrix square root, inverse square root, or other central functions.

- Prove a direct instance optimality bound for the matrix exponential. Some progress in [Druskin, Greenbaum, Knizhnerman '98].
- Prove instance optimality bounds for the matrix square root, inverse square root, or other central functions.
- Understand the role of finite precision. We know that it matters a lot: uniform approximation bounds are much more stable than instance optimal ones.

Theorem (Musco, Musco, Sidford, 2018)

For any bounded function f, if Lanczos is run on a finite precision computer with $\log(\text{poly}(n, \kappa, \delta_k))$ *bits of precision,*

$$\|\mathbf{Q}f(\mathbf{T})\mathbf{Q}^T\mathbf{b} - f(\mathbf{A})\mathbf{b}\| \leq 7k \cdot \delta_k \cdot \|\mathbf{b}\|,$$

where

$$\delta_k = \min_{\text{degree } k \text{ poly } p} \left(\max_{x \in [\lambda_{\min}(\mathbf{A}), \lambda_{\max}(\mathbf{A})]} |f(x) - p(x)| \right).$$

I.e., the uniform approximation bound basically goes through with a small additional constant factor.

The story is much more complicated for near-optimality bounds, and we know relative error guarantees do not hold.

In particular, there is always a degree n polynomial with zero error in approximating f at $\mathbf{A}$'s eigenvalues. So a finite-precision near optimality bound would e.g. imply that $\mathbf{A}^{-1}\mathbf{b}$ can be computed in:

$$O(\text{nnz}(\mathbf{A})n) \text{ time },$$

independent of the condition number.

In finite precision, Lanczos/CG do no achieve this, but there is some really cool recent progress on faster solvers for sparse systems using Krylov methods [Peng, Vempala '21, Nie '22].

Greenbaum (1989): Finite precision Lanczos and conjugate gradient match the best polynomial approximating $1/x$ in tiny intervals around A's eigenvalues:

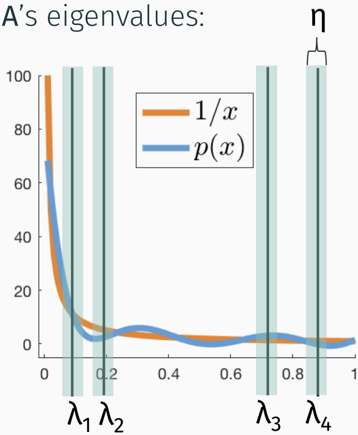

THANK YOU!