Instance Optimal Iterative Methods for Matrix Function Approximation

Cameron Musco

University of Massachusetts Amherst

Joint with: Noah Amsel (NYU), Tyler Chen (NYU), Anne Greenbaum (Univ Washington), Christopher Musco (NYU)

Basic problem:

- Consider a scalar function $f : \mathbb{R} \to \mathbb{R}$.
- For symmetric $A \in \mathbb{R}^{n\times n}$ with eigendecomposition $A = \sum_{i=1}^n \lambda_i v_i v_i^T$, define the matrix function $f(A) = \sum_{i=1}^n f(\lambda_i) v_i v_i^T$.

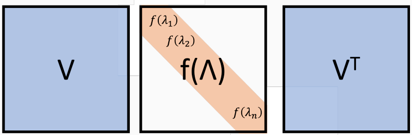

Matrix Function Approximation

Basic problem:

- Consider a scalar function $f : \mathbb{R} \to \mathbb{R}$.
- For symmetric $A \in \mathbb{R}^{n\times n}$ with eigendecomposition $A = \sum_{i=1}^n \lambda_i v_i v_i^T$, define the matrix function $f(A) = \sum_{i=1}^n f(\lambda_i) v_i v_i^T$.

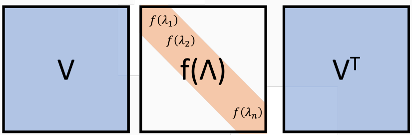

- Given $b \in \mathbb{R}^n$ we would like to compute the matrix-vector product $f(A)b$.
- For general A, 'exact' computation requires $O(n^\omega)$ time (i.e., roughly a full eigendecomposition).
- We will thus seek approximation algorithms that are much faster.

Example Applications

- When $f(x) = 1/x$, $f(A) = A^{-1}$ and $A^{-1}b$ is the solution to a linear system.
- When A is PSD (i.e., has non-negative eigenvalues) and $f(x) = \sqrt{x}$, $f(A) = A^{1/2}$ is the matrix squareroot. Needed e.g., to sample from a multivariate Gaussian distribution with covariance A.
- In many cases, the trace of $f(A)$ is of interest since $\text{tr}(f(A)) = \sum_{i=1}^{n} f(\lambda_i)$. E.g., when $f(x) = log(x)$, $\text{tr}(f(A)) = \sum_{i=1}^{n} \log(\lambda_i) = \text{logdet}(A)$.
- $\text{tr}(f(A))$ can be estimated by repeatedly multiplying $f(A)$ by random b (Hutchinson's method).

Example Applications

- When $f(x) = 1/x$, $f(A) = A^{-1}$ and $A^{-1}b$ is the solution to a linear system.
- When A is PSD (i.e., has non-negative eigenvalues) and $f(x) = \sqrt{x}$, $f(A) = A^{1/2}$ is the matrix squareroot. Needed e.g., to sample from a multivariate Gaussian distribution with covariance A.
- In many cases, the trace of $f(A)$ is of interest since $\text{tr}(f(A)) = \sum_{i=1}^{n} f(\lambda_i)$. E.g., when $f(x) = log(x)$, $\text{tr}(f(A)) = \sum_{i=1}^{n} \log(\lambda_i) = \text{logdet}(A)$.
- $\text{tr}(f(A))$ can be estimated by repeatedly multiplying $f(A)$ by random b (Hutchinson's method).

Other important matrix functions: The matrix sign function, step functions, the matrix exponential, the matrix squareroot.

Krylov Subspace Methods

Krylov subspace methods are the dominant approach to approximating matrix functions.

- Key idea: when $f(x)$ is a degree-q polynomial, $f(A)$ can be computed with just q matrix-vector products with A. At most $O(n^2 \cdot q)$ run time – faster for sparse or structured A.

$$f(A)b = c_0 b + c_1 Ab + c_2 A^2 b + \ldots + c_q A^q b.$$

- Krylov methods approximate general functions as polynomials, and then use the above fact to compute them quickly.

Krylov Subspace Methods

Krylov subspace methods are the dominant approach to approximating matrix functions.

- Key idea: when $f(x)$ is a degree-q polynomial, $f(A)$ can be computed with just q matrix-vector products with A. At most $O(n^2 \cdot q)$ run time – faster for sparse or structured A.

$$f(A)b = c_0 b + c_1 A b + c_2 A^2 b + \ldots + c_q A^q b.$$

- Krylov methods approximate general functions as polynomials, and then use the above fact to compute them quickly.
- In this talk we will focus on the Lanczos method, which can be used to approximate any $f(A)$ and is very popular in practice.
- If A is PSD and $f(x) = 1/x$, the Lanczos method gives the same output as the conjugate gradient (CG) algorithm.

Krylov Subspace Methods

Krylov subspace methods are the dominant approach to approximating matrix functions.

- Key idea: when $f(x)$ is a degree-q polynomial, $f(A)$ can be computed with just q matrix-vector products with A. At most $O(n^2 \cdot q)$ run time – faster for sparse or structured A.

$$f(A)b = c_0 b + c_1 Ab + c_2 A^2 b + \ldots + c_q A^q b.$$

- Krylov methods approximate general functions as polynomials, and then use the above fact to compute them quickly.
- In this talk we will focus on the Lanczos method, which can be used to approximate any $f(A)$ and is very popular in practice.
- If A is PSD and $f(x) = 1/x$, the Lanczos method gives the same output as the conjugate gradient (CG) algorithm.

Other examples: MINRES, gradient descent, accelerated gradient descent, and many other iterative methods for linear systems.

The Lanczos Method

- The Lanczos method run for k iterations employs $k-1$ matrix vector products with A and computes $Q \in \mathbb{R}^{n \times k}$ with orthonormal columns that span the Krylov subspace $\{b, Ab, A^2b, \ldots, A^{k-1}b\}$.
- The method orthogonalizes Q via a tri-term recurrence which ensures that $T = Q^T AQ$ is tridiagonal.
- We then approximate $f(A)b \approx Qf(T)Q^T b$.

The Lanczos Method

- The Lanczos method run for k iterations employs $k-1$ matrix vector products with A and computes $Q \in \mathbb{R}^{n \times k}$ with orthonormal columns that span the Krylov subspace $\{b, Ab, A^2b, \ldots, A^{k-1}b\}$.
- The method orthogonalizes Q via a tri-term recurrence which ensures that $T = Q^T AQ$ is tridiagonal.
- We then approximate $f(A)b \approx Qf(T)Q^T b$.
- $f(T)$ can be computed in $O(k^2) = O(nk)$ time, so the total runtime of the method is just $k \cdot mvm(A) + O(nk)$.

The Lanczos Method

- The Lanczos method run for k iterations employs $k-1$ matrix vector products with A and computes $Q \in \mathbb{R}^{n \times k}$ with orthonormal columns that span the Krylov subspace $\{b, Ab, A^2b, \ldots, A^{k-1}b\}$.
- The method orthogonalizes Q via a tri-term recurrence which ensures that $T = Q^T AQ$ is tridiagonal.
- We then approximate $f(A)b \approx Qf(T)Q^T b$.
- $f(T)$ can be computed in $O(k^2) = O(nk)$ time, so the total runtime of the method is just $k \cdot mvm(A) + O(nk)$.
- Observe that for $i < k$, $A^i b = QQ^T A^i b = QQ^T A^i QQ^T b = QT^i Q^T b$.

The Lanczos Method

- The Lanczos method run for k iterations employs $k-1$ matrix vector products with A and computes $Q \in \mathbb{R}^{n \times k}$ with orthonormal columns that span the Krylov subspace $\{b, Ab, A^2b, \ldots, A^{k-1}b\}$.
- The method orthogonalizes Q via a tri-term recurrence which ensures that $T = Q^T AQ$ is tridiagonal.
- We then approximate $f(A)b \approx Qf(T)Q^T b$.
- $f(T)$ can be computed in $O(k^2) = O(nk)$ time, so the total runtime of the method is just $k \cdot mvm(A) + O(nk)$.
- Observe that for $i < k$, $A^i b = QQ^T A^i b = QQ^T A^i QQ^T b = QT^i Q^T b$.
- So, by linearity, if p is a polynomial of degree $< k$, the method is exact. I.e., $p(A)b = Qp(T)Q^T b$.

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$\|f(A)b - Qf(T)Q^T b\|_2 \leq \|p(A)b - Qp(T)Q^T b\|_2 \\ + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^T b - Qp(T)Q^T b\|_2$$

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$\begin{aligned}\|f(A)b - Qf(T)Q^T b\|_2 &\leq \|p(A)b - Qp(T)Q^T b\|_2 \\ &\quad + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^T b - Qp(T)Q^T b\|_2 \\ &\leq \|f(A) - p(A)\|_2 + \|f(T) - p(T)\|_2\end{aligned}$$

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$\begin{aligned}
\|f(A)b - Qf(T)Q^T b\|_2 &\leq \|p(A)b - Qp(T)Q^T b\|_2 \\
&\qquad + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^T b - Qp(T)Q^T b\|_2 \\
&\leq \|f(A) - p(A)\|_2 + \|f(T) - p(T)\|_2 \\
&\leq \max_{\lambda_i(A)} |f(\lambda_i) - p(\lambda_i)| + \max_{\lambda_i(T)} |f(\lambda_i) - p(\lambda_i)|
\end{aligned}$$

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$
\begin{aligned}
\|f(A)b - Qf(T)Q^T b\|_2 &\leq \|p(A)b - Qp(T)Q^T b\|_2 \\
&\qquad + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^T b - Qp(T)Q^T b\|_2 \\
&\leq \|f(A) - p(A)\|_2 + \|f(T) - p(T)\|_2 \\
&\leq \max_{\lambda_i(A)} |f(\lambda_i) - p(\lambda_i)| + \max_{\lambda_i(T)} |f(\lambda_i) - p(\lambda_i)| \\
&\leq 2 \cdot \max_{\lambda \in [\lambda_{min}(A), \lambda_{max}(A)]} |f(\lambda) - p(\lambda)|.
\end{aligned}
$$

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$\begin{aligned}
\|f(A)b - Qf(T)Q^T b\|_2 &\leq \|p(A)b - Qp(T)Q^T b\|_2 \\
&\qquad + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^T b - Qp(T)Q^T b\|_2 \\
&\leq \|f(A) - p(A)\|_2 + \|f(T) - p(T)\|_2 \\
&\leq \max_{\lambda_i(A)} |f(\lambda_i) - p(\lambda_i)| + \max_{\lambda_i(T)} |f(\lambda_i) - p(\lambda_i)| \\
&\leq 2 \cdot \max_{\lambda \in [\lambda_{min}(A), \lambda_{max}(A)]} |f(\lambda) - p(\lambda)|.
\end{aligned}$$

The above holds for any polynomial p. By optimizing over p we have:

$$\|f(A)b - Qf(T)Q^T b\|_2 \leq 2 \cdot \min_{\{p:\ \text{degree p < k}\}} \max_{\lambda \in [\lambda_{min}(A), \lambda_{max}(A)]} |f(\lambda) - p(\lambda)|.$$

Uniform Error Bound for the Lanczos Method

Via triangle inequality, we get a really nice error bound for approximating general matrix functions.

Assume $\|b\|_2 = 1$. Then for any polynomial p with degree $< k$,

$$\begin{aligned}
\|f(A)b - Qf(T)Q^Tb\|_2 &\leq \|p(A)b - Qp(T)Q^Tb\|_2 \\
&\quad + \|f(A)b - p(A)b\|_2 + \|Qf(T)Q^Tb - Qp(T)Q^Tb\|_2 \\
&\leq \|f(A) - p(A)\|_2 + \|f(T) - p(T)\|_2 \\
&\leq \max_{\lambda_i(A)} |f(\lambda_i) - p(\lambda_i)| + \max_{\lambda_i(T)} |f(\lambda_i) - p(\lambda_i)| \\
&\leq 2 \cdot \max_{\lambda \in [\lambda_{min}(A), \lambda_{max}(A)]} |f(\lambda) - p(\lambda)|.
\end{aligned}$$

The above holds for any polynomial p. By optimizing over p we have:

$$\|f(A)b - Qf(T)Q^Tb\|_2 \leq 2 \cdot \min_{\{p:\text{ degree p < k}\}} \max_{\lambda \in [\lambda_{min}(A), \lambda_{max}(A)]} |f(\lambda) - p(\lambda)|.$$

I.e., Lanczos gives within a two factor of the best uniform approximation error of f by a polynomial on A's spectral range.

Uniform Error Bound for the Lanczos Method

- The uniform convergence bound for Lanczos is very powerful.
- It can be used e.g. to show that CG solves linear systems to accuracy ϵ in $O(\sqrt{\kappa(A)} \cdot \log 1/\epsilon)$ iterations.
- It is robust to roundoff error [Druskin, Knizhnerman '91], [Musco, Musco, Sidford '18].
- It can be shown to be tight up to a factor 2 for any continuous f and worst case A, b, even when $n = k + 1$.

Uniform Error Bound for the Lanczos Method

- The uniform convergence bound for Lanczos is very powerful.
- It can be used e.g. to show that CG solves linear systems to accuracy ϵ in $O(\sqrt{\kappa(A)} \cdot \log 1/\epsilon)$ iterations.
- It is robust to roundoff error [Druskin, Knizhnerman '91], [Musco, Musco, Sidford '18].
- It can be shown to be tight up to a factor 2 for any continuous f and worst case A, b, even when $n = k + 1$.
- But the uniform approximation bound almost always fails to capture the very strong performance of Lanczos in practice.
- This gap between theory and practice is what our work seeks to address.

Performance in Practice

In practice, Lanczos often far outperforms the uniform error bound. It is often within a small constant factor of the best approximation in the Krylov subspace. I.e., of $\min_{\{p:\ \text{degree } p < k\}} \|f(A)b - p(A)b\|_2$.

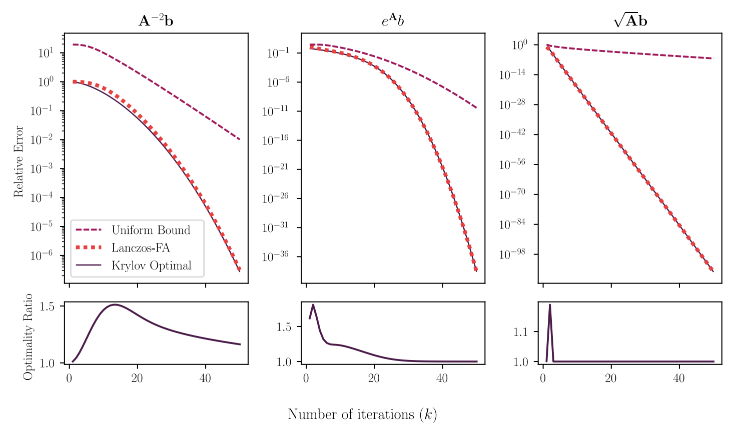

Instance Optimality Bounds

Our Goal: Show that for common matrix functions f,

$$\|f(A)b - Qf(T)Q^T b\|_2 \leq C \cdot \min_{\{p:\text{ degree p < k}\}} \|f(A)b - p(A)b\|_2,$$

for some reasonably small approximation factor C.

- Note that this 'instance optimality guarantee' is always at most the uniform approximation bound, and often is smaller by a wide margin.

Instance Optimality Bounds

Our Goal: Show that for common matrix functions f,

$$\|f(A)b - Qf(T)Q^T b\|_2 \leq C \cdot \min_{\{p:\text{ degree p < k}\}} \|f(A)b - p(A)b\|_2,$$

for some reasonably small approximation factor C.

- Note that this 'instance optimality guarantee' is always at most the uniform approximation bound, and often is smaller by a wide margin.
- When $f(x) = 1/x$ (the linear system case), Lanczos is instance optimal for $C = \sqrt{\kappa(A)}$.
- A related guarantee is was shown for the matrix exponential by [Druskin, Greenbaum, Knizhnerman '98].
- But we are not aware of any other known results for important functions like the matrix sign function, square root, etc.

Our Result: Instance Optimality Bounds for Rational Functions

Setting:

- Let $r(x) = \frac{p(x)}{(x-z_1)(x-z_2)\ldots(x-z_q)}$ be a degree-(m, q) rational function with real poles lying outside the spectral range of A. I.e., $z_1, \ldots, z_q \notin [\lambda_{min}(A), \lambda_{max}(A)]$.
- Let $A_i = A - z_i I$.

Our Result: Instance Optimality Bounds for Rational Functions

Setting:

- Let $r(x) = \frac{p(x)}{(x-z_1)(x-z_2)\ldots(x-z_q)}$ be a degree-(m, q) rational function with real poles lying outside the spectral range of A. I.e., $z_1, \ldots, z_q \notin [\lambda_{min}(A), \lambda_{max}(A)]$.
- Let $A_i = A - z_i I$.

Main Theorem: Lanczos is instance optimal for a such a rational function with $C = q \cdot \prod_{i=1}^{q} \kappa(A_i)$. Specifically, for $k \geq \max\{m, q-1\}$,

$$\|f(A)b - Qf(T)Q^T b\|_2 \leq C \cdot \min_{\{p:\text{ degree p < k-q+1}\}} \|f(A)b - p(A)b\|_2.$$

Remarks on the Main Result

- Rational functions are interesting in their own right. They include e.g. $1/x$, $1/x^q$, etc.
- More importantly, they often give very accurate approximations to functions with discontinuities, like the squareroot or step functions.
- Our error bound can be used to give stronger error bounds for Lanczos in approximating such functions.
- Our approximation factor $C = q \cdot \prod_{i=1}^{q} \kappa(A_i)$ is really bad. Grows exponentially in q. We believe it can be significantly improved!
- The best empirical lower bound we observe for C when all poles are at 0 is roughly $\sqrt{q \cdot \kappa(A)}$.

Empirical Performance

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.

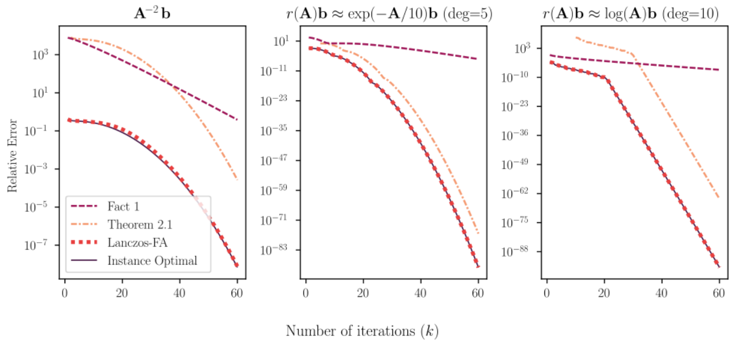

Empirical Performance

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.

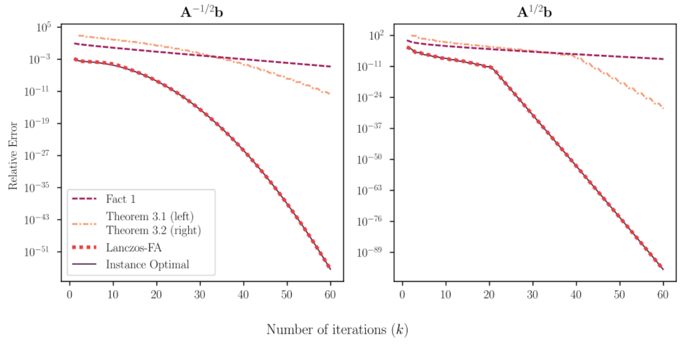

Proof Sketch for the Inverse

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

Proof Sketch for the Inverse

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

$$\|A^{-1}b - QT^{-1}Q^Tb\|_2 \leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2$$

Proof Sketch for the Inverse

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

$$\begin{aligned}\|A^{-1}b - QT^{-1}Q^Tb\|_2 &\leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2 \\ &= \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}Q(Q^TAQ)^{-1}Q^TA^{1/2}A^{-1/2}b\|_2\end{aligned}$$

Proof Sketch for the Inverse

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

$$\begin{aligned}
\|A^{-1}b - QT^{-1}Q^Tb\|_2 &\leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}Q(Q^TAQ)^{-1}Q^TA^{1/2}A^{-1/2}b\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1/2}b - A^{1/2}Qx\|_2
\end{aligned}$$

Proof Sketch for the Inverse

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

$$\begin{aligned}
\|A^{-1}b - QT^{-1}Q^Tb\|_2 &\leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}Q(Q^TAQ)^{-1}Q^TA^{1/2}A^{-1/2}b\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1/2}b - A^{1/2}Qx\|_2 \\
&\leq \sqrt{\frac{\lambda_{max}(A)}{\lambda_{min}(A)}} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1}b - Qx\|_2
\end{aligned}$$

Proof Sketch for the Inverse

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

$$
\begin{aligned}
\|A^{-1}b - QT^{-1}Q^Tb\|_2 &\leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}Q(Q^TAQ)^{-1}Q^TA^{1/2}A^{-1/2}b\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1/2}b - A^{1/2}Qx\|_2 \\
&\leq \sqrt{\frac{\lambda_{max}(A)}{\lambda_{min}(A)}} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1}b - Qx\|_2 \\
&= \sqrt{\kappa(A)} \cdot \min_{\{p:\text{ degree p < k}\}} \|A^{-1}b - p(A)b\|_2
\end{aligned}
$$

Proof Sketch for the Inverse

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

$$\begin{aligned}
\|A^{-1}b - QT^{-1}Q^Tb\|_2 &\leq \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}QT^{-1}Q^Tb\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \|A^{-1/2}b - A^{1/2}Q(Q^TAQ)^{-1}Q^TA^{1/2}A^{-1/2}b\|_2 \\
&= \sqrt{\lambda_{max}(A)} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1/2}b - A^{1/2}Qx\|_2 \\
&\leq \sqrt{\frac{\lambda_{max}(A)}{\lambda_{min}(A)}} \cdot \min_{x \in \mathbb{R}^k} \|A^{-1}b - Qx\|_2 \\
&= \sqrt{\kappa(A)} \cdot \min_{\{p\text{: degree p < k}\}} \|A^{-1}b - p(A)b\|_2
\end{aligned}$$

Another view: Lanczos computes the A-norm optimal approximation to $A^{-1}b$ in the Krylov subspace. This is within a $\sqrt{\kappa(A)}$ factor of the best ℓ_2 norm approximation.

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^T b\|_2 = \|A^{-2}b - QT^{-1}Q^T A^{-1}b\|_2 + \|QT^{-2}Q^T b - QT^{-1}Q^T A^{-1}b\|_2.$$

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2.$$

Term 1: $\|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2$.

- This is the best approximation to $A^{-2}b$ in the span of the Krylov subspace in the A-norm. Following the same proof as in the $f(x) = 1/x$ case we have:

$$\|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 \leq \sqrt{\kappa(A)} \cdot \min_{\{p:\text{ degree p < k}\}} \|A^{-2}b - p(A)b\|_2.$$

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$.

Term 2: $\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2.$$

Term 2: $\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$

$$\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 \leq \|QT^{-1}Q^T\|_2 \cdot \|QT^{-1}Q^Tb - A^{-1}b\|_2$$

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2.$$

Term 2: $\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$

$$\begin{aligned}
\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 &\leq \|QT^{-1}Q^T\|_2 \cdot \|QT^{-1}Q^Tb - A^{-1}b\|_2 \\
&\leq \frac{\sqrt{\kappa(A)}}{\lambda_{min}(A)} \cdot \min_{\{p:\ \text{degree p < k}\}} \|A^{-1}b - p(A)b\|_2.
\end{aligned}$$

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2.$$

Term 2: $\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$

$$\begin{aligned}\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 &\leq \|QT^{-1}Q^T\|_2 \cdot \|QT^{-1}Q^Tb - A^{-1}b\|_2 \\ &\leq \frac{\sqrt{\kappa(A)}}{\lambda_{min}(A)} \cdot \min_{\{p:\ \text{degree p < k}\}} \|A^{-1}b - p(A)b\|_2.\end{aligned}$$

Key Idea: The optimal error for approximating A^{-1} with degree k can be bounded by the optimal error for approximating A^{-2} with degree $k-1$. Since $\|A^{-1}b - p(A)Ab\|_2 \leq \lambda_{max}(A) \cdot \|A^{-2}b - p(A)b\|_2$.

Proof Sketch

To get the main idea of how to generalize this to rational functions, it suffices to consider the special case of $r(x) = 1/x^2$. I.e., $r(A) = A^{-2}$.

$$\|A^{-2}b - QT^{-2}Q^Tb\|_2 = \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2.$$

Term 2: $\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2$

$$\begin{aligned}\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 &\leq \|QT^{-1}Q^T\|_2 \cdot \|QT^{-1}Q^Tb - A^{-1}b\|_2 \\ &\leq \frac{\sqrt{\kappa(A)}}{\lambda_{min}(A)} \cdot \min_{\{p:\text{ degree p < k}\}} \|A^{-1}b - p(A)b\|_2.\end{aligned}$$

Key Idea: The optimal error for approximating A^{-1} with degree k can be bounded by the optimal error for approximating A^{-2} with degree $k-1$. Since $\|A^{-1}b - p(A)Ab\|_2 \leq \lambda_{max}(A) \cdot \|A^{-2}b - p(A)b\|_2$.

Overall, this gives:

$$\|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 \leq \kappa(A)^{3/2} \cdot \min_{\{p:\text{ degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2.$$

Proof Sketch

Overall, we have:

$$
\begin{aligned}
\|A^{-2}b - QT^{-2}Q^Tb\|_2 &= \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 \\
&\leq \sqrt{\kappa(A)} \cdot \min_{\{p\text{: degree p < k}\}} \|A^{-2}b - p(A)b\|_2 \\
&\qquad + \kappa(A)^{3/2} \cdot \min_{\{p\text{: degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2 \\
&\leq 2\kappa(A)^{3/2} \cdot \min_{\{p\text{: degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2.
\end{aligned}
$$

Proof Sketch

Overall, we have:

$$\begin{aligned}\|A^{-2}b - QT^{-2}Q^Tb\|_2 &= \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 \\ &\leq \sqrt{\kappa(A)} \cdot \min_{\{p\text{: degree p < k}\}} \|A^{-2}b - p(A)b\|_2 \\ &\quad + \kappa(A)^{3/2} \cdot \min_{\{p\text{: degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2 \\ &\leq 2\kappa(A)^{3/2} \cdot \min_{\{p\text{: degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2.\end{aligned}$$

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

Proof Sketch

Overall, we have:

$$\begin{aligned}
\|A^{-2}b - QT^{-2}Q^Tb\|_2 &= \|A^{-2}b - QT^{-1}Q^TA^{-1}b\|_2 + \|QT^{-2}Q^Tb - QT^{-1}Q^TA^{-1}b\|_2 \\
&\leq \sqrt{\kappa(A)} \cdot \min_{\{p:\text{ degree p < k}\}} \|A^{-2}b - p(A)b\|_2 \\
&\qquad + \kappa(A)^{3/2} \cdot \min_{\{p:\text{ degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2 \\
&\leq 2\kappa(A)^{3/2} \cdot \min_{\{p:\text{ degree p < k-1}\}} \|A^{-2}b - p(A)b\|_2.
\end{aligned}$$

- 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, or show stronger lower bounds. Our best numerical lower bound for A^{-q} is $C = \sqrt{q\kappa}$, as compared to our best theoretical upper bound of $C = q\kappa^q$.

Open Questions

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

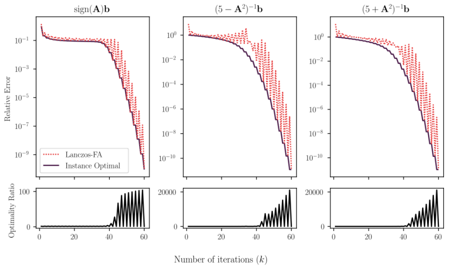

Open Questions

- Tighten our bounds, or show stronger lower bounds. Our best numerical lower bound for A^{-q} is $C = \sqrt{q\kappa}$, as compared to our best theoretical upper bound of $C = q\kappa^q$.
- Extend our results to the case when $r(x)$ has poles in A's spectral range. In this case, Lanczos seems to be oscillate between very bad and near optimal solutions.
- Can explain when A is not PSD and $r(x) = 1/x$ by relating the convergence of CG to that of MINRES. But lack a general result.
- Prove an instance optimality bound for the matrix exponential. Some progress in [Druskin, Greenbaum, Knizhnerman '98].
- Prove instance optimality bounds for the matrix squareroot, inverse squareroot, or other central functions.

Open Questions

- Tighten our bounds, or show stronger lower bounds. Our best numerical lower bound for A^{-q} is $C = \sqrt{q\kappa}$, as compared to our best theoretical upper bound of $C = q\kappa^q$.
- Extend our results to the case when $r(x)$ has poles in A's spectral range. In this case, Lanczos seems to be oscillate between very bad and near optimal solutions.
- Can explain when A is not PSD and $r(x) = 1/x$ by relating the convergence of CG to that of MINRES. But lack a general result.
- Prove an instance optimality bound for the matrix exponential. Some progress in [Druskin, Greenbaum, Knizhnerman '98].
- Prove instance optimality bounds for the matrix squareroot, inverse squareroot, 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.