

Low Rank Approximation

Lecture 1

Daniel Kressner

Chair for Numerical Algorithms and HPC

Institute of Mathematics, EPFL

`daniel.kressner@epfl.ch`




ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE



Organizational aspects

- ▶ **Lecture dates:** 16.4., 23.4., 30.4., 14.5., 28.5., 4.6., 11.6., 18.6., 25.6., 2.7. (tentative)
- ▶ **Exam:** To be discussed next week (most likely oral exam).
- ▶ **Webpage:** https://www5.in.tum.de/wiki/index.php/Low_Rank_Approximation
Slides on <http://anchp.epfl.ch>.
- ▶ **EFY = Exercise For You.**



This is how Cambridge Analytica's Facebook targeting model really worked — according to the person who built it

The method was similar to the one Netflix uses to recommend movies — no crystal ball, but good enough to make an effective political tool.

By **MATTHEW HINDMAN** March 30, 2018, 11:35 a.m.



People read news differently (i.e., worse) on phones than they do on desktop, new research suggests

Laura Hazard Owen



... his [Aleksandr Kogan's] message went on to confirm that his approach was indeed similar to *SVD or other matrix factorization* methods, like in the Netflix Prize competition, and the Kosinski-Stillwell-Graepel Facebook model. *Dimensionality reduction* of Facebook data was the core of his model.

Rank and matrix factorizations

For field F , let $A \in F^{m \times n}$. Then

$$\text{rank}(A) := \dim(\text{range}(A)).$$

For simplicity, $F = \mathbb{R}$ throughout the lecture and often $m \geq n$.

Let $\mathcal{B} = \{b_1, \dots, b_r\} \subset \mathbb{R}^m$ with $r = \text{rank}(A)$ be basis of $\text{range}(A)$.

Then each of the columns of $A = (a_1, a_2, \dots, a_n)$ can be expressed as linear combination of \mathcal{B} :

$$a_j = \sum_{i=1}^r b_i c_{ij} \quad \text{for some coefficients } c_{ij} \in \mathbb{R}, \quad i = 1, \dots, r, \quad j = 1, \dots, n.$$

Defining $B = (b_1, b_2, \dots, b_r) \in \mathbb{R}^{m \times r}$:

$$a_j = B \begin{pmatrix} c_{j1} \\ \vdots \\ c_{jr} \end{pmatrix} \rightsquigarrow A = B \begin{pmatrix} c_{11} & \cdots & c_{n1} \\ \vdots & & \vdots \\ c_{1r} & \cdots & c_{nr} \end{pmatrix}$$


Rank and matrix factorizations

Lemma. A matrix $A \in \mathbb{R}^{m \times n}$ of rank r admits a factorization of the form

$$A = BC^T, \quad B \in \mathbb{R}^{m \times r}, \quad C \in \mathbb{R}^{n \times r}.$$

We say that A has **low rank** if $\text{rank}(A) \ll m, n$.

Illustration of low-rank factorization:



	A	BC^T
#entries	mn	$mr + nr$

- ▶ Generically (and in most applications), A has **full rank**, that is, $\text{rank}(A) = \min\{m, n\}$.
- ▶ Aim instead at **approximating** A by a low-rank matrix.

Questions addressed in lecture series

- What?** Theoretical foundations of low-rank approximation.
- When?** A priori and a posteriori estimates for low-rank approximation. Situations that allow for low-rank approximation techniques.
- Why?** Applications in engineering, scientific computing, data analysis, ... where low-rank approximation plays a central role.
- How?** State-of-the-art algorithms for performing and working with low-rank approximations.

Will cover both, matrices and tensors.

Contents of Lecture 1

1. Fundamental tools (SVD, relation to eigenvalues, norms, best low-rank approximation)
2. Overview of applications
3. Fundamental tools (Stability, QR)
4. Extensions (weighted approximation, bivariate functions)
5. Subspace iteration

Literature for Lecture 1

Golub/Van Loan'2013 Golub, Gene H.; Van Loan, Charles F. Matrix computations. Fourth edition. Johns Hopkins University Press, Baltimore, MD, 2013.

Horn/Johnson'2013 Horn, Roger A.; Johnson, Charles R. Matrix analysis. Second edition. Cambridge University Press, 2013.

+ References on slides.

1. Fundamental tools

- ▶ SVD
- ▶ Relation to eigenvalues
- ▶ Norms
- ▶ Best low-rank approximation

The singular value decomposition

Theorem (SVD). Let $A \in \mathbb{R}^{m \times n}$ with $m \geq n$. Then there are orthogonal matrices $U \in \mathbb{R}^{m \times m}$ and $V \in \mathbb{R}^{n \times n}$ such that

$$A = U\Sigma V^T, \quad \text{with} \quad \Sigma = \begin{pmatrix} \sigma_1 & & & \\ & \ddots & & \\ & & \sigma_n & \\ & & & 0 \end{pmatrix} \in \mathbb{R}^{m \times n}$$

and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$.

- ▶ $\sigma_1, \dots, \sigma_n$ are called singular values
- ▶ u_1, \dots, u_m are called *left* singular vectors
- ▶ v_1, \dots, v_n are called *right* singular vectors
- ▶ $Av_i = \sigma_i u_i$, $A^T u_i = \sigma_i v_i$ for $i = 1, \dots, n$.
- ▶ Singular values are always uniquely defined by A .
- ▶ Singular values are *never* unique. If $\sigma_1 > \sigma_2 > \dots > \sigma_n > 0$ then unique up to $u_i \leftarrow \pm u_i$, $v_i \leftarrow \pm v_i$.

SVD: Sketch of proof

Induction over n . $n = 1$ trivial.

For general n , let v_1 solve $\max\{\|Av\|_2 : \|v\|_2 = 1\} =: \|A\|_2$. Set $\sigma_1 := \|A\|_2$ and $u_1 := Av_1/\sigma_1$.¹ By definition,

$$Av_1 = \sigma_1 u_1.$$

After completion to orthogonal matrices $U_1 = (u_1, U_\perp) \in \mathbb{R}^{m \times m}$ and $V_1 = (v_1, V_\perp) \in \mathbb{R}^{n \times n}$:

$$U_1^T AV_1 = \left(\begin{array}{c|c} u_1^T Av_1 & u_1^T AV_\perp \\ \hline U_\perp^T Av_1 & U_\perp^T AV_\perp \end{array} \right) = \left(\begin{array}{c|c} \sigma_1 & w^T \\ \hline 0 & A_1 \end{array} \right),$$

with $w := V_\perp^T A^T u_1$ and $A_1 = U_\perp^T AV_\perp$. $\|\cdot\|_2$ invariant under orthogonal transformations \rightsquigarrow

$$\sigma_1 = \|A\|_2 = \|U_1^T AV_1\|_2 = \left\| \left(\begin{array}{c|c} \sigma_1 & w^T \\ \hline 0 & A_1 \end{array} \right) \right\|_2 \geq \sqrt{\sigma_1^2 + \|w\|_2^2}.$$

Hence, $w = 0$. Proof completed by applying induction to A_1 .

¹If $\sigma_1 = 0$, choose arbitrary u_1 .

Very basic properties of the SVD

- ▶ $r = \text{rank}(A)$ is number of nonzero singular values of A .
- ▶ $\text{kernel}(A) = \text{span}\{v_{r+1}, \dots, v_n\}$
- ▶ $\text{range}(A) = \text{span}\{u_1, \dots, u_r\}$

SVD: Computation (for small dense matrices)

Computation of SVD proceeds in two steps:

1. **Reduction to bidiagonal form:** By applying n Householder reflectors from left and $n - 1$ Householder reflectors from right, compute orthogonal matrices U_1, V_1 such that

$$U_1^T A V_1 = B = \begin{pmatrix} B_1 \\ 0 \end{pmatrix} = \begin{pmatrix} \diagdown \\ 0 \end{pmatrix},$$

that is, $B_1 \in \mathbb{R}^{n \times n}$ is an upper bidiagonal matrix.

2. **Reduction to diagonal form:** Use Divide&Conquer to compute orthogonal matrices U_2, V_2 such that $\Sigma = U_2^T B_1 V_2$ is diagonal. Set $U = U_1 U_2$ and $V = V_1 V_2$.

Step 1 is usually the most expensive. Remarks on Step 1:

- ▶ If m is significantly larger than n , say, $m \geq 3n/2$, first computing QR decomposition of A reduces cost.
- ▶ Most modern implementations reduce A successively via banded form to bidiagonal form.²

²Bischof, C. H.; Lang, B.; Sun, X. A framework for symmetric band reduction. ACM Trans. Math. Software 26 (2000), no. 4, 581–601.

SVD: Computation (for small dense matrices)

In most applications, vectors u_{n+1}, \dots, u_m are not of interest. By omitting these vectors one obtains the following variant of the SVD.

Theorem (Economy size SVD). Let $A \in \mathbb{R}^{m \times n}$ with $m \geq n$. Then there is a matrix $U \in \mathbb{R}^{m \times n}$ with orthonormal columns and an orthonormal matrix $V \in \mathbb{R}^{n \times n}$ such that

$$A = U\Sigma V^T, \quad \text{with} \quad \Sigma = \begin{pmatrix} \sigma_1 & & \\ & \ddots & \\ & & \sigma_n \end{pmatrix} \in \mathbb{R}^{n \times n}$$

and $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n \geq 0$.

Computed by MATLAB's `[U, S, V] = svd(A, 'econ')`.


Complexity:

	memory	operations
singular values only	$\mathcal{O}(mn)$	$\mathcal{O}(mn^2)$
economy size SVD	$\mathcal{O}(mn)$	$\mathcal{O}(mn^2)$
(full) SVD	$\mathcal{O}(m^2 + mn)$	$\mathcal{O}(m^2n + mn^2)$

SVD: Computation (for small dense matrices)

Beware of roundoff error when interpreting singular value plots.

Example: `semilogy(svd(hilb(100)))`



- ▶ Kink is caused by roundoff error and does not reflect true behavior of singular values.
- ▶ Exact singular values are known to decay exponentially.³
- ▶ *Sometimes* more accuracy possible.⁴

³Beckermann, B. The condition number of real Vandermonde, Krylov and positive definite Hankel matrices. Numer. Math. 85 (2000), no. 4, 553–577.

⁴Drmač, Z.; Veselić, K. New fast and accurate Jacobi SVD algorithm. I. SIAM J. Matrix Anal. Appl. 29 (2007), no. 4, 1322–1342

Singular/eigenvalue relations: symmetric matrices

Symmetric $A = A^T \in \mathbb{R}^{n \times n}$ admits **spectral decomposition**

$$A = U \operatorname{diag}(\lambda_1, \lambda_2, \dots, \lambda_n) U^T$$

with orthogonal matrix U .

After reordering may assume $|\lambda_1| \geq |\lambda_2| \geq \dots \geq |\lambda_n|$. Spectral decomposition can be turned into SVD $A = U \Sigma V^T$ by defining

$$\Sigma = \operatorname{diag}(|\lambda_1|, \dots, |\lambda_n|), \quad V = U \operatorname{diag}(\operatorname{sign}(\lambda_1), \dots, \operatorname{sign}(\lambda_n)).$$

Remark: This extends to the more general case of normal matrices (e.g., orthogonal or symmetric) via complex spectral or real Schur decompositions.

Singular/eigenvalue relations: general matrices

Consider SVD $A = U\Sigma V^T$ of $A \in \mathbb{R}^{m \times n}$ with $m \geq n$. We then have:

1. Spectral decomposition of Gramian

$$A^T A = V \Sigma^T \Sigma V^T = V \operatorname{diag}(\sigma_1^2, \dots, \sigma_n^2) V^T \rightsquigarrow$$

$A^T A$ has eigenvalues $\sigma_1^2, \dots, \sigma_n^2$,

right singular vectors of A are eigenvectors of $A^T A$.

2. Spectral decomposition of Gramian

$$A A^T = U \Sigma \Sigma^T U^T = U \operatorname{diag}(\sigma_1^2, \dots, \sigma_n^2, 0, \dots, 0) U^T \rightsquigarrow$$

$A A^T$ has eigenvalues $\sigma_1^2, \dots, \sigma_n^2$ and, additionally, $m - n$ zero eigenvalues,

first n left singular vectors A are eigenvectors of $A A^T$.

3. Decomposition of Golub-Kahan matrix

$$\mathcal{A} = \begin{pmatrix} 0 & A \\ A^T & 0 \end{pmatrix} = \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} 0 & \Sigma \\ \Sigma^T & 0 \end{pmatrix} \begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix}^T.$$

EFY. Prove that \mathcal{A} has eigenvalues $\pm \sigma_j$ with eigenvectors $\frac{1}{\sqrt{2}} \begin{pmatrix} \pm u_j \\ v_j \end{pmatrix}$.

Norms: Spectral and Frobenius norm

Given SVD $A = U\Sigma V^T$, one defines:

- ▶ Spectral norm: $\|A\|_2 = \sigma_1$.
- ▶ Frobenius norm: $\|A\|_F = \sqrt{\sigma_1^2 + \cdots + \sigma_n^2}$.

Basic properties:

- ▶ $\|A\|_2 = \max\{\|Av\|_2 : \|v\|_2 = 1\}$ (see proof of SVD).
- ▶ $\|\cdot\|_2$ and $\|\cdot\|_F$ are both (submultiplicative) matrix norms.
- ▶ $\|\cdot\|_2$ and $\|\cdot\|_F$ are both unitarily invariant, that is

$$\|QAZ\|_2 = \|A\|_2, \quad \|QAZ\|_F = \|A\|_F$$

for any orthogonal matrices Q, Z .

- ▶ $\|A\|_2 \leq \|A\|_F \leq \|A\|_2 / \sqrt{r}$
- ▶ $\|AB\|_F \leq \min\{\|A\|_2 \|B\|_F, \|A\|_F \|B\|_2\}$

EFY. Prove these two inequalities. Hint for the second inequality: Use the relations on the next slide to first show that $\|B\|_F = \|(\|b_1\|_2, \dots, \|b_n\|_2)\|_F$.

EFY. Find a matrix $A \in \mathbb{R}^{m_1 \times n}$ and a nonzero matrix $B \in \mathbb{R}^{m_2 \times n}$ such that $\|A\|_2 = \left\| \begin{pmatrix} A \\ B \end{pmatrix} \right\|_2$. Classify the set of matrices

$A \in \mathbb{R}^{m_1 \times n}$ such that $\|A\|_2 < \left\| \begin{pmatrix} A \\ B \end{pmatrix} \right\|_2$ for every nonzero matrix $B \in \mathbb{R}^{m_2 \times n}$.

Investigate analogous questions for the Frobenius norm.

Euclidean geometry on matrices

Let $B \in \mathbb{R}^{n \times n}$ have eigenvalues $\lambda_1, \dots, \lambda_n \in \mathbb{C}$. Then

$$\text{trace}(B) := b_{11} + \dots + b_{nn} = \lambda_1 + \dots + \lambda_n.$$

In turn,

$$\|A\|_F^2 = \text{trace } A^T A = \text{trace } A A^T = \sum_{i,j} a_{ij}^2.$$

Two simple consequences:

- ▶ $\|\cdot\|_F$ is the norm induced by the **matrix inner product**

$$\langle A, B \rangle := \text{trace}(A B^T), \quad A, B \in \mathbb{R}^{m \times n}.$$

- ▶ Partition $A = (a_1, a_2, \dots, a_n)$ and define **vectorization**

$$\text{vec}(A) = \begin{pmatrix} a_1 \\ \vdots \\ a_n \end{pmatrix} \in \mathbb{R}^{mn}.$$

Then $\langle A, B \rangle = \langle \text{vec}(A), \text{vec}(B) \rangle$ and $\|A\|_F = \|\text{vec}(A)\|_2$.

Von Neumann's trace inequality

Theorem

For $m \geq n$, let $A, B \in \mathbb{R}^{m \times n}$ have singular values $\sigma_1(A) \geq \dots \geq \sigma_n(A)$ and $\sigma_1(B) \geq \dots \geq \sigma_n(B)$, respectively. Then

$$|\langle A, B \rangle| \leq \sigma_1(A)\sigma_1(B) + \dots + \sigma_n(A)\sigma_n(B).$$

Consequence:

$$\begin{aligned} \|A - B\|_F^2 &= \langle A - B, A - B \rangle = \|A\|_F^2 - 2\langle A, B \rangle + \|B\|_F^2 \\ &\geq \|A\|_F^2 - 2 \sum_{i=1}^n \sigma_i(A)\sigma_i(B) + \|B\|_F^2 \\ &= \sum_{i=1}^n (\sigma_i(A) - \sigma_i(B))^2. \end{aligned}$$

EFY. Use Von Neumann's trace inequality and the SVD to show for $1 \leq k \leq n$ that

$$\max\{|\langle A, PQ^T \rangle| : P \in \mathbb{R}^{m \times k}, Q \in \mathbb{R}^{n \times k}, P^T P = Q^T Q = I_k\} = \sigma_1(A) + \dots + \sigma_k(A).$$

Proof of Von Neumann's trace inequality⁵

Singular value vector $\sigma(A)$ can be written as convex combination

$$\sigma(A) = \sigma_n(A)f_n + (\sigma_{n-1}(A) - \sigma_n(A))f_{n-1} + \dots$$

with $f_j = e_1 + \dots + e_j$. Decompose A analogously via its SVD

$$A = U_A \Sigma_A V_A^T:$$

$$A = \sigma_n(A)A_n + (\sigma_{n-1}(A) - \sigma_n(A))A_{n-1} + \dots, \quad A_j := U_A \text{diag}(f_j) V_A^T$$

Insert in lhs of trace inequality:

$$|\langle A, B \rangle| \leq \sigma_n(A)|\langle A_n, B \rangle| + (\sigma_{n-1}(A) - \sigma_n(A))|\langle A_{n-1}, B \rangle| + \dots$$

Rhs is linear wrt $\sigma(A) \rightsquigarrow$ May assume $A = A_k$ for $k = 1, \dots, n$.

Analogously for B .

⁵This proof follows [Grigorieff, R. D. Note on von Neumann's trace inequality. Math. Nachr. 151 (1991), 327–328]. For Mirsky's ingenious proof based on doubly stochastic matrices; see Theorem 8.7.6 in [Horn/Johnson'2013].

Proof of Von Neumann's trace inequality

Let $A = U_A \text{diag}(f_k) V_A^T$, $B = U_B \text{diag}(f_\ell) V_B^T$, and $k \leq \ell$. Then

$$\begin{aligned}\langle A, B \rangle &= \text{trace} \left(\sum_{i=1}^k v_{A,i} u_{A,i}^T \sum_{j=1}^{\ell} u_{B,j} v_{B,j}^T \right) \\ &= \sum_{i=1}^k \sum_{j=1}^{\ell} \text{trace} (v_{A,i} u_{A,i}^T u_{B,j} v_{B,j}^T) \\ &= \sum_{i=1}^k \sum_{j=1}^{\ell} (u_{A,i}^T u_{B,j}) (v_{B,j}^T v_{A,i})\end{aligned}$$

Cauchy-Schwartz \rightsquigarrow

$$|\langle A, B \rangle| \leq \sum_{i=1}^k \|U_B^T u_{A,i}\|_2 \|V_B^T v_{A,i}\|_2 = k,$$

which completes the proof.

Schatten norms

There are other unitarily invariant matrix norms.⁶

Let $s(A) = (\sigma_1, \dots, \sigma_n)$. The p -Schatten norm defined by

$$\|A\|_{(p)} := \|s(A)\|_p$$

is a matrix norm for any $1 \leq p \leq \infty$.

$p = \infty$: spectral norm, $p = 2$: Frobenius norm, $p = 1$: nuclear norm.

EFY. What is $\lim_{p \rightarrow 0^+} \|A\|_{(p)}$?

Definition

The dual of a matrix norm $\|\cdot\|$ on $\mathbb{R}^{m \times n}$ is defined by

$$\|A\|^D = \max\{\langle A, B \rangle : \|B\| = 1\}.$$

Lemma

Let $p, q \in [1, \infty]$ such that $p^{-1} + q^{-1} = 1$. Then

$$\|A\|_{(p)}^D = \|A\|_{(q)}.$$

EFY. Prove this lemma for $p = \infty$. Hint: Von Neumann's trace inequality.

⁶Complete characterization via symm gauge functions in [Horn/Johnson'2013].

Best low-rank approximation

Consider $k < n$ and let

$$U_k := (u_1 \ \cdots \ u_k), \quad \Sigma_k := \text{diag}(\sigma_1, \dots, \sigma_k), \quad V_k := (v_1 \ \cdots \ v_k).$$

Then

$$\mathcal{T}_k(A) := U_k \Sigma_k V_k^T$$

has rank at most k . For any unitarily invariant norm $\|\cdot\|$:

$$\|\mathcal{T}_k(A) - A\| = \|\text{diag}(0, \dots, 0, \sigma_{k+1}, \dots, \sigma_n)\|$$

In particular, for spectral norm and the Frobenius norm:

$$\|A - \mathcal{T}_k(A)\|_2 = \sigma_{k+1}, \quad \|A - \mathcal{T}_k(A)\|_F = \sqrt{\sigma_{k+1}^2 + \cdots + \sigma_n^2}.$$

Nearly equal if and only if singular values decay sufficiently quickly.

Best low-rank approximation

Theorem (Schmidt-Mirsky). Let $A \in \mathbb{R}^{m \times n}$. Then

$$\|A - \mathcal{T}_k(A)\| = \min \{ \|A - B\| : B \in \mathbb{R}^{m \times n} \text{ has rank at most } k \}$$

holds for any unitarily invariant norm $\| \cdot \|$.

Proof⁷ for $\| \cdot \|_F$: Follows directly from consequence of Von Neumann's trace inequality.

Proof for $\| \cdot \|_2$: For any $B \in \mathbb{R}^{m \times n}$ of rank $\leq k$, $\text{kernel}(B)$ has dimension $\geq n - k$. Hence, $\exists w \in \text{kernel}(B) \cap \text{range}(V_{k+1})$ with $\|w\|_2 = 1$. Then

$$\begin{aligned} \|A - B\|_2^2 &\geq \|(A - B)w\|_2^2 = \|Aw\|_2^2 = \|AV_{k+1}V_{k+1}^T w\|_2^2 \\ &= \|U_{k+1}\Sigma_{k+1}V_{k+1}^T w\|_2^2 \\ &= \sum_{j=1}^{r+1} \sigma_j |v_j^T w|^2 \geq \sigma_{k+1} \sum_{j=1}^{r+1} |v_j^T w|^2 = \sigma_{k+1}. \end{aligned}$$

⁷See Section 7.4.9 in [Horn/Johnson'2013] for the general case.

Best low-rank approximation

Uniqueness:

- ▶ If $\sigma_k > \sigma_{k+1}$ best rank- k approximation with respect to Frobenius norm is unique.
- ▶ If $\sigma_k = \sigma_{k+1}$ best rank- k approximation never unique. For example I_3 has several best rank-two approximations:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

- ▶ With respect to spectral norm best rank- k approximation only unique if $\sigma_{k+1} = 0$. For example, $\text{diag}(2, 1, \epsilon)$ with $0 < \epsilon < 1$ has infinitely many best rank-two approximations:

$$\begin{pmatrix} 2 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 2 - \epsilon/2 & 0 & 0 \\ 0 & 1 - \epsilon/2 & 0 \\ 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 2 - \epsilon/3 & 0 & 0 \\ 0 & 1 - \epsilon/3 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \dots$$

EFY. Given a symmetric matrix $A \in \mathbb{R}^{n \times n}$ and $1 \leq k < n$, show that there is always a best rank- k approximation that is symmetric. Is every best rank- k approximation (with respect to Frobenius norm) symmetric? What about the spectral norm?

Approximating the range of a matrix

Aim at finding a matrix $Q \in \mathbb{R}^{m \times k}$ with orthonormal columns such that

$$\text{range}(Q) \approx \text{range}(A).$$

$I - QQ^T$ is orthogonal projector onto $\text{range}(Q)^\perp \rightsquigarrow$

Aim at minimizing

$$\|(I - QQ^T)A\| = \|A - QQ^T A\|$$

for unitarily invariant norm $\|\cdot\|$. Because $\text{rank}(QQ^T A) \leq k$,

$$\|A - QQ^T A\| \geq \|A - \mathcal{T}_k(A)\|.$$

Setting $Q = U_k$ one obtains

$$U_k U_k^T A = U_k U_k^T U \Sigma V^T = U_k \Sigma_k V_k^T = \mathcal{T}_k(A).$$

$\rightsquigarrow Q = U_k$ is optimal.

Approximating the range of a matrix

Variation:

$$\max\{\|Q^T A\|_F : Q^T Q = I_k\}.$$

Equivalent to

$$\max\{|\langle AA^T, QQ^T \rangle| : Q^T Q = I_k\}.$$

By Von Neumann's trace inequality and equivalence between eigenvectors of AA^T and left singular vectors of A , optimal Q given by U_k .

EFY. When replacing the Frobenius norm by the spectral norm in this formulation, does one obtain the same result?


2. Applications

- ▶ Principal Component Analysis
- ▶ Matrix Completion
- ▶ Some other applications

Principal Component Analysis (PCA)

- ▶ Most popular method for **dimensionality reduction** in statistics, data science, ...

Consider N independently drawn observations for K random variables X_1, \dots, X_K . Illustration of $N = 100$ observations for $K = 2$:




Principal Component Analysis (PCA)

Each of the observations is arranged in a vector $x_j \in \mathbb{R}^K$ with $j = 1, \dots, N$.

Subtract sample mean

$$\bar{x} := \frac{1}{N}(x_1 + \dots + x_N)$$

Data with mean subtracted:



Principal Component Analysis (PCA)

Covariance matrix

$$C := \frac{1}{N-1} \sum_{j=1}^N (x_j - \bar{x})(x_j - \bar{x})^T.$$

Diagonal entry c_{ii} estimates variance of X_i , while off-diagonal entry c_{ik} estimates covariance between X_i and X_k .

Defining $A := [x_1 - \bar{x}, \dots, x_N - \bar{x}] \in \mathbb{R}^{K \times N}$, we can equivalently write

$$C = \frac{1}{N-1} AA^T.$$

Principal Component Analysis (PCA)

Reduce data to dimension 1:

Find linear combination $Y_1 = w_1 X_1 + \dots + w_K X_K$ with $w_1, \dots, w_K \in \mathbb{R}$ and $w_1^2 + \dots + w_K^2 = 1$ that captures most of the observed variation.

Maximize variance of new variable Y_1 .

Corresponding observations of Y_1 given by $w^T x_1, \dots, w^T x_N$ with sample mean $w^T \bar{x} \rightsquigarrow$ maximization of variance corresponds to


$$\max_{\substack{w \in \mathbb{R}^N \\ \|w\|_2=1}} \sum_{j=1}^N (w^T x_j - w^T \bar{x})^2 = \max_{\substack{w \in \mathbb{R}^N \\ \|w\|_2=1}} \|w^T A\|_2^2.$$

Optimal vector w given by dominant left singular vector of $A!$
(Corresponds to eigenvector for largest eigenvalue of AA^T .)


This is the **first principal vector**.

Principal Component Analysis (PCA)

Data with first principal vector:



Projection of data onto first principal vector:



Principal Component Analysis (PCA)

- ▶ Analogously, first k principal vectors given by dominant k left singular vectors u_1, \dots, u_k . Equivalent to best rank- k approximation of data matrix:

$$\min_{U^T U = I_k} \|A - UC^T\|$$

- ▶ PCA **not** robust wrt outliers in data. Robust PCA⁸ uses model

$$A \approx \text{low rank} + \text{sparse}.$$

Obtained via solution of

$$\min\{\|L\|_{(1)} + \lambda\|S\|_1 : A = L + S\},$$

for multiplier $\lambda > 0$ and $\|S\|_1 = 1$ -norm of $\text{vec}(S)$.

⁸Emmanuel J. Candes; Xiaodong Li; Yi Ma; John Wright. Robust Principal Component Analysis?

Matrix Completion

				
Alice	4			4
Bob		5	4	
Joe		5		
Sam	5			

Assume that data matrix is modeled by (low) rank k . Two popular approaches to deal with missing entries:

1. Impute data (insert 0 or row/column means in missing entries). Apply SVD to get best low-rank approximation BC^T of imputed data matrix.
2. Find rank- k matrix BC^T that fits known entries best; measured in (weighted) Euclidean norm.

Predict unknown entries from BC^T .

Netflix prize won by combination of matrix completion with other techniques.

Applications in Scientific Computing and Engineering

- ▶ POD, reduced basis method, reduced-order modelling.
- ▶ High-dimensional integration.
- ▶ Solution of large-scale matrix equations. Optimal control.
- ▶ Solution of high-dimensional PDEs.
- ▶ Uncertainty quantification.
- ▶ ...

Several of these will be covered in later parts of the course.

3. Fundamental Tools

- ▶ Stability of SVD
- ▶ Canonical angles
- ▶ Stability of low-rank approximation
- ▶ QR decomposition

Stability of SVD

What happens to SVD if A is perturbed by noise?

Lemma

Let $A, E \in \mathbb{R}^{m \times n}$. Then

$$|\sigma_i(A + E) - \sigma_i(A)| \leq \|E\|_2.$$

Proof.

Using the characterization

$$\sigma_i(A + E) = \min\{\|B\|_2 : \text{rank}(A + E - B) \leq i - 1\}$$

and setting $B = A - \mathcal{T}_{i-1}(A) + E$, we obtain

$$\sigma_i(A + E) \leq \|B\|_2 \leq \|A - \mathcal{T}_{i-1}(A)\|_2 + \|E\|_2 = \sigma_i(A) + \|E\|_2,$$

which implies the result. □

Result also special case of famous Weyl's inequality.

EFY. Show that the matrix rank is a lower semi-continuous function.

Stability of SVD

Singular values are perfectly well conditioned.

Singular vectors tend to be less stable! Example:

$$A = \begin{pmatrix} 1 & 0 \\ 0 & 1 + \varepsilon \end{pmatrix}, \quad E = \begin{pmatrix} 0 & \varepsilon \\ \varepsilon & -\varepsilon \end{pmatrix}.$$

- ▶ A has right singular vectors $\begin{pmatrix} 1 \\ 0 \end{pmatrix}, \begin{pmatrix} 0 \\ 1 \end{pmatrix}$.
- ▶ $A + E$ has right singular vectors $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$

To formulate perturbation bound, need to measure distances between subspaces.

Canonical angles

Let columns of $X, Y \in \mathbb{C}^{n \times k}$ contain orthonormal bases of two k -dimensional subspaces $\mathcal{X}, \mathcal{Y} \subset \mathbb{C}^n$, respectively. Denote singular values (in reverse order) of $X^T Y$:

$$0 \leq \sigma_1 \leq \dots \leq \sigma_p \leq 1.$$

We call

$$\theta_i(\mathcal{X}, \mathcal{Y}) := \arccos \sigma_i, \quad i = 1, \dots, k,$$

the **canonical angles between \mathcal{X} and \mathcal{Y}** . Note: For $k = 1$, θ_1 is the usual angle $\theta(x, y)$ between vectors.

Geometric characterization:

$$\theta_1(\mathcal{X}, \mathcal{Y}) = \max_{\substack{x \in \mathcal{X} \\ x \neq 0}} \min_{\substack{y \in \mathcal{Y} \\ y \neq 0}} \theta(x, y).$$

It follows that $\theta_1(\mathcal{X}, \mathcal{Y}) = 0$ if and only if $\mathcal{X} \cap \mathcal{Y}^\perp \neq \{0\}$.

Canonical angles

Note that XX^T and YY^T are orthogonal projectors on \mathcal{X} and \mathcal{Y} , respectively.

Lemma (Projector characterization)

Define $\sin \Theta(\mathcal{X}, \mathcal{Y}) = \text{diag}(\sin \theta_1(\mathcal{X}, \mathcal{Y}), \dots, \sin \theta_p(\mathcal{X}, \mathcal{Y}))$. Then

$$\sin \theta_1(\mathcal{X}, \mathcal{Y}) = \|\sin \Theta(\mathcal{X}, \mathcal{Y})\|_2 = \|XX^T - YY^T\|_2.$$

Proof. See Theorem I.5.5 in [Stewart/Sun'1990].

Lemma

Let $Q \in \mathbb{R}^{(n-k) \times k}$, and $\mathcal{X} = \text{range} \begin{pmatrix} I_k \\ 0 \end{pmatrix}$, $\mathcal{Y} = \text{range} \begin{pmatrix} I_p \\ Q \end{pmatrix}$.

Then $\theta_1(\mathcal{X}, \mathcal{Y}) = \arctan \|Q\|_2$.

Proof.

The columns of $\begin{pmatrix} I_p \\ Q \end{pmatrix} (I + Q^T Q)^{-1/2}$ form an orthonormal basis of \mathcal{Y} . By definition, this implies that $\cos \theta_1(\mathcal{X}, \mathcal{Y})$ is the smallest singular value of $\begin{pmatrix} I + Q^T Q \end{pmatrix}^{-1/2}$. By the SVD of Q , it follows that

$$\cos \theta_1(\mathcal{X}, \mathcal{Y}) = \frac{1}{\sqrt{1 + \|Q\|_2^2}}.$$

Stability of SVD

Theorem (Wedin). Let $k < n$ and assume

$$\delta := \sigma_k(A + E) - \sigma_{k+1}(A) > 0.$$

Let $\mathcal{U}_k/\tilde{\mathcal{U}}_k/\mathcal{V}_k/\tilde{\mathcal{V}}_k$ denote subspaces spanned by first k left/right singular vectors of $A / A + E$. Then

$$\sqrt{\|\sin \Theta(\mathcal{U}_k, \tilde{\mathcal{U}}_k)\|_F^2 + \|\sin \Theta(\mathcal{V}_k, \tilde{\mathcal{V}}_k)\|_F^2} \leq \sqrt{2} \frac{\|E\|_F}{\delta}. \quad (1)$$

Θ : diagonal matrix containing canonical angles between two subspaces.

- ▶ Perturbation on input multiplied by $\delta^{-1} \approx [\sigma_k(A) - \sigma_{k+1}(A)]^{-1}$.
- ▶ Bad news for stability of low-rank approximations?

Stability of low-rank approximation

Lemma. Let $A \in \mathbb{R}^{m \times n}$ have rank $\leq k$. Then

$$\|\mathcal{T}_k(A + E) - A\| \leq C\|E\|$$

holds with $C = 2$ for any unitarily invariant norm $\|\cdot\|$. For the Frobenius norm, the constant can be improved to $C = (1 + \sqrt{5})/2$.

Proof. Schmidt-Mirsky gives $\|\mathcal{T}_k(A + E) - (A + E)\| \leq \|E\|$. Triangle inequality implies

$$\|\mathcal{T}_k(A + E) - (A + E) + (A + E) - A\| \leq 2\|E\|.$$

Second part is result by Hackbusch⁹. □

Implication for **general** matrix A :

$$\begin{aligned}\|\mathcal{T}_k(A + E) - \mathcal{T}_k(A)\| &= \|\mathcal{T}_k(\mathcal{T}_k(A) + (A - \mathcal{T}_k(A)) + E) - \mathcal{T}_k(A)\| \\ &\leq C\|(A - \mathcal{T}_k(A)) + E\| \leq C(\|A - \mathcal{T}_k(A)\| + \|E\|).\end{aligned}$$

Perturbations on the level of truncation error pose no danger.

⁹Hackbusch, W. New estimates for the recursive low-rank truncation of block-structured matrices. Numer. Math. 132 (2016), no. 2, 303–328

Stability of low-rank approximation: Application

Consider partitioned matrix

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}, \quad A_{ij} \in \mathbb{R}^{m_i \times n_j},$$

and desired rank $k \leq m_i, n_j$. Let $\varepsilon := \|\mathcal{T}_k(A) - A\|$.

$$E_{ij} := \mathcal{T}_k(A_{ij}) - A_{ij} \quad \Rightarrow \quad \|E_{ij}\| \leq \varepsilon.$$

By stability of low-rank approximation,

$$\left\| \mathcal{T}_k \begin{pmatrix} \mathcal{T}_k(A_{11}) & \mathcal{T}_k(A_{12}) \\ \mathcal{T}_k(A_{21}) & \mathcal{T}_k(A_{22}) \end{pmatrix} - A \right\|_F = \left\| \mathcal{T}_k \left(A + \begin{pmatrix} E_{11} & E_{12} \\ E_{21} & E_{22} \end{pmatrix} \right) - A \right\|_F \leq C\varepsilon,$$

with $C = \frac{3}{2}(1 + \sqrt{5})$.

This allows, e.g., to perform truncations in parallel.

The QR decomposition

Theorem

Let $X \in \mathbb{R}^{m \times n}$ with $m \geq n$. Then there is an orthogonal matrix $Q \in \mathbb{R}^{m \times m}$ such that

$$X = QR, \quad \text{with} \quad R = \begin{pmatrix} R_1 \\ 0 \end{pmatrix} = \begin{pmatrix} \triangle \\ 0 \end{pmatrix},$$

that is, $R_1 \in \mathbb{R}^{n \times n}$ is an upper triangular matrix.

MATLAB: `[Q,R] = qr(X)`.

Will use **economy size QR decomposition** instead: Letting $Q_1 \in \mathbb{C}^{m \times n}$ contain first n columns of Q , one obtains

$$X = Q_1 R_1 = Q_1 \cdot \triangle.$$

MATLAB: `[Q,R] = qr(X,0)`.

EFY. Let $A = (a_1, a_2, \dots, a_n)$ with $a_j \in \mathbb{R}^m$. Using the QR decomposition, show *Hadamard's inequality*:

$$|\det(A)| \leq \|a_1\|_2 \cdot \|a_2\|_2 \cdots \|a_n\|_2.$$

Characterize the set of all $m \times n$ matrices A for which equality holds.

QR for recompression

Suppose that

$$A = BC^T, \quad \text{with } B \in \mathbb{R}^{m \times K}, C \in \mathbb{R}^{n \times K}. \quad (2)$$

Goal: Compute best rank- k approximation of A for $k < K$.

Typical example: **Sum** of J matrices of rank k :

$$A = \sum_{j=1}^J \underbrace{B_j}_{\in \mathbb{R}^{m \times k}} \underbrace{C_j^T}_{\in \mathbb{R}^{n \times k}} = \underbrace{(B_1 \ \cdots \ B_J)}_{\mathbb{R}^{m \times Jk}} \underbrace{(C_1 \ \cdots \ C_J)^T}_{\mathbb{R}^{m \times Jk}}. \quad (3)$$

Algorithm to recompress A :

1. Compute (economic) QR decomps $B = Q_B R_B$ and $C = Q_C R_C$.
2. Compute truncated SVD $\mathcal{T}_k(R_B R_C^T) = \tilde{U}_k \Sigma_k \tilde{V}_k^T$.
3. Set $U_k = Q_B \tilde{U}_k$, $V_k = Q_C \tilde{V}_k$ and return $\mathcal{T}_k(A) := U_k \Sigma_k V_k^T$.

Returns best rank- k approximation of A with $O((m+n)K^2)$ ops.

4. Extensions

- ▶ Weighted and structured low-rank approximation
- ▶ Semi-separable approximation of bivariate functions

Weighted low-rank approximation

If some columns or rows are more important than others (e.g., they are known to be less corrupted by noise), replace low-rank approximation problem by

$$\min \{ \|D_R(A - B)D_C\| : B \in \mathbb{R}^{m \times n} \text{ has rank at most } k \}$$

with suitably chosen pos def diagonal matrices D_R, D_C . More general:

Given invertible matrices $W_R \in \mathbb{R}^{m \times m}$, $W_C \in \mathbb{R}^{n \times n}$, **weighted low-rank approximation** problem consists of

$$\min \{ \|W_R(A - B)W_C\| : B \in \mathbb{R}^{m \times n} \text{ has rank at most } k \}.$$

Solution given by

$$B = W_R^{-1} \cdot \mathcal{T}_k(W_R A W_C) \cdot W_C^{-1}$$

Proof: EFY.

Remark: Numerically more stable approach via generalized SVD [Golub/Van Loan'2013].

Limit case: Infinite weights

Choosing diagonal weights that converge to $\infty \rightsquigarrow$ rows/columns remain unperturbed.

Case of fixed columns: Consider block column partition

$$A = \begin{pmatrix} A_1 & A_2 \end{pmatrix}.$$

Consider

$$\min \{ \|A_2 - B_2\| : \begin{pmatrix} A_1 & B_2 \end{pmatrix} \text{ has rank at most } k \}$$

No/trivial solution if $\text{rank}(A_1) \geq k$. Assume $\ell := \text{rank}(A_1) < k$ and let $X_1 \in \mathbb{R}^{n \times \ell}$ contain orthonormal basis of $\text{range}(A_1)$. Then¹⁰

$$B_2 = X_1 X_1^T A_2 + \mathcal{T}_{k-\ell}((I - X_1 X_1^T) A_2).$$

¹⁰Golub, G. H.; Hoffman, A.; Stewart, G. W. A generalization of the Eckart-Young-Mirsky matrix approximation theorem. *Linear Algebra Appl.* 88/89 (1987), 317–327.

General weights

Given an $mn \times mn$ symmetric pos def matrix W , define

$$\|A\|_W = \sqrt{\text{vec}(A)^T W \text{vec}(A)}$$

Equals Frobenius norm for $W = I$. General weighted low-rank approximation problem:

$$\min \{ \|A - B\|_W : B \in \mathbb{R}^{m \times n} \text{ has rank at most } k \}.$$

EFY. Show that this problem can be rephrased as the previously considered (standard) weighted low-rank approximation problem for the case of a Kronecker product $W = W_2 \otimes W_1$. Hint: Cholesky decomposition.

- ▶ For general W no expression in terms of SVD available \rightsquigarrow need to use general optimization method.
- ▶ Similarly, imposing general structures on A (such as nonnegativity, fixing individual entries, ...) usually does not admit solutions in terms of SVD. Often end up with NP-hard problems.

Separable approximation of bivariate functions

Given $\Omega_x \subset \mathbb{R}^{d_x}$ and $\Omega_y \subset \mathbb{R}^{d_y}$ aim at finding semi-separable approximation of $f \in L^2(\Omega_x \times \Omega_y) \cong L^2(\Omega_x) \otimes L^2(\Omega_y)$:

$$f(x, y) \approx g_1(x)h_1(y) + \cdots + g_r(x)h_r(y)$$

for $g_1, \dots, g_r \in L^2(\Omega_x)$, $h_1, \dots, h_r \in L^2(\Omega_y)$

Application to higher-dimensional integrals:

$$\begin{aligned} & \int_{\Omega_x} \int_{\Omega_y} f(x, y) d\mu_y(y) d\mu_x(x) \\ & \approx \sum_{i=1}^r \int_{\Omega_x} \int_{\Omega_y} g_i(x)h_i(y) d\mu_y(y) d\mu_x(x) \\ & = \sum_{i=1}^r \left[\int_{\Omega_x} g_i(x) d\mu_x(x) \right] \left[\int_{\Omega_y} h_i(y) d\mu_y(y) \right] \end{aligned}$$

\rightsquigarrow semi-separable approximation breaks down dimensionality of integrals (for separable measures).

Separable approximation of bivariate functions

Given $f \in L^2(\Omega_x \times \Omega_y)$, consider linear operator

$$L_f : L^2(\Omega_x) \rightarrow L^2(\Omega_y), \quad w \mapsto \int_{\Omega_x} w(x)f(x, y) dx.$$

Admits SVD

$$L_f(\cdot) = \sum_{i=1}^{\infty} \sigma_i u_i \langle v_i, \cdot \rangle$$

with L^2 orthonormal bases u_1, u_2, \dots and v_1, v_2, \dots

Best semi-separable approximation of f (in $L^2(\Omega_x \times \Omega_y)$) given by

$$f_r(x, y) = \sum_{i=1}^r \sigma_i u_i(x) v_i(y),$$

provided that $\sum_{i=1}^{\infty} \sigma_i^2 < \infty$ (Hilbert-Schmidt).

$$\|f - f_r\|_{L^2}^2 = \sigma_{r+1}^2 + \sigma_{r+2}^2 + \dots$$

Separable and low-rank approximation

Choose discretization $x_1, \dots, x_m \in \Omega_x$, $y_1, \dots, y_n \in \Omega_y$. Define

$$F = \begin{pmatrix} f(x_1, y_1) & f(x_1, y_2) & \cdots & f(x_1, y_n) \\ f(x_2, y_1) & f(x_2, y_2) & \cdots & f(x_2, y_n) \\ \vdots & \vdots & \ddots & \vdots \\ f(x_m, y_1) & f(x_m, y_2) & \cdots & f(x_m, y_n) \end{pmatrix}$$

and

$$F_r = \begin{pmatrix} f_r(x_1, y_1) & f_r(x_1, y_2) & \cdots & f_r(x_1, y_n) \\ f_r(x_2, y_1) & f_r(x_2, y_2) & \cdots & f_r(x_2, y_n) \\ \vdots & \vdots & \ddots & \vdots \\ f_r(x_m, y_1) & f_r(x_m, y_2) & \cdots & f_r(x_m, y_n) \end{pmatrix} = \sum_{i=1}^r \begin{pmatrix} g_i(x_1) \\ g_i(x_2) \\ \vdots \\ g_i(x_m) \end{pmatrix} \begin{pmatrix} h_i(y_1) \\ h_i(y_2) \\ \vdots \\ h_i(y_n) \end{pmatrix}^T$$

F_r has rank at most r .

EFY. Prove $\|F - F_r\|_F^2 \leq \sigma_{r+1}^2 + \sigma_{r+2}^2 + \cdots$.

5. Subspace Iteration

Subspace iteration and low-rank approximation

Subspace iteration = extension of power method.

Input: Matrix $A \in \mathbb{R}^{m \times n}$.

- 1: Choose starting matrix $X^{(0)} \in \mathbb{R}^{m \times k}$ with $(X^{(0)})^T X^{(0)} = I_k$.
- 2: $j = 0$.
- 3: **repeat**
- 4: Set $j := j + 1$.
- 5: Compute $Y^{(j)} := AA^T X^{(j-1)}$.
- 6: Compute economy size QR factorization: $Y^{(j)} = QR$.
- 7: Set $X^{(j)} := Q$.
- 8: **until** convergence is detected

As will soon be seen, converges to basis of dominant subspace \mathcal{U}_k .

Low-rank approximation obtained from

$$\mathcal{T}_k(A) \approx X^{(j)} (X^{(j)})^T A.$$

Convergence of subspace iteration

Theorem

Consider SVD $A = U\Sigma V^T$ and, for $k < n$, let $\mathcal{U}_k = \text{span}\{u_1, \dots, u_k\}$. Assume that $\sigma_{k+1} > \sigma_k$ and $\theta_1(\mathcal{U}_k, \mathcal{X}^{(0)}) < \pi/2$. Then the iterates $\mathcal{X}^{(j)} = \text{range}(X^{(j)})$ of the subspace iteration satisfy

$$\tan \theta_1(\mathcal{U}_k, \mathcal{X}^{(j)}) = \left| \frac{\sigma_{k+1}}{\sigma_k} \right|^{2j} \tan \theta_1(\mathcal{U}_k, \mathcal{X}^{(0)}).$$

Sketch of proof. As angles do not depend on choice of bases, may omit QR decompositions $\rightsquigarrow X^{(j)} = (AA^T)^j X^{(0)}$. By SVD of A , may set $A = \Sigma$ and hence $\mathcal{U}_k = \text{span}\{e_1, \dots, e_k\}$. Partition

$$\Sigma = \begin{pmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{pmatrix}, \quad X^{(0)} = \begin{pmatrix} X_1^{(0)} \\ X_2^{(0)} \end{pmatrix}, \quad \Sigma_1, X_1^{(0)} \in \mathbb{R}^{k \times k}.$$


Result follows from applying expression for tangent of θ_1 . □

EFY. Complete details of proof.

Numerical experiments

Convergence of subspace iteration for 100×100 Hilbert matrix.

$k = 5 \rightsquigarrow \sigma_{k+1}/\sigma_k = 0.188$. Random starting guess.



Black curve:

$$\tan \theta_1(\mathcal{U}_k, \mathcal{X}^{(j)})$$

Blue curve:

$$\|\mathcal{T}_7(\mathbf{A}) - \mathbf{X}^{(j)}(\mathbf{X}^{(j)})^T \mathbf{A}\|_2$$

Red curve:


$$\|\mathbf{A} - \mathbf{X}^{(j)}(\mathbf{X}^{(j)})^T \mathbf{A}\|_2$$

Numerical experiments

Convergence of subspace iteration for matrix with singular values

$$1, 0.99, 0.98, \frac{1}{10}, \frac{0.99}{10}, \frac{0.98}{10}, \frac{1}{100}, \frac{0.99}{100}, \frac{0.98}{100}, \dots$$

$k = 7 \rightsquigarrow \sigma_{k+1}/\sigma_k = 0.99$. Random starting guess.



Black curve:

$$\tan \theta_1(\mathcal{U}_k, \mathcal{X}^{(j)})$$

Blue curve:

$$\|\mathcal{T}_7(\mathbf{A}) - \mathbf{X}^{(j)}(\mathbf{X}^{(j)})^T \mathbf{A}\|_2$$

Red curve:

$$\|\mathbf{A} - \mathbf{X}^{(j)}(\mathbf{X}^{(j)})^T \mathbf{A}\|_2$$

Numerical experiments

Observations:

- ▶ Low-rank approximation sufficiently good (for most purposes) already after 1 iteration.
- ▶ Convergence to dominant subspace arbitrarily slow, but not relevant.
- ▶ Classical, asymptotic convergence analysis insufficient.
- ▶ Pre-asymptotic analysis needs to take randomization of starting guess into account.