# Randomized Algorithms in Numerical Linear Algebra

Yuji Nakatsukasa
Oxford University

Based on joint work with
Maike Meier (Oxford), Joel Tropp (Caltech)
references: arXiv 2111.00113, 2009.11392, 2105.07388

*JSIAM Applied Maths Seminar 2022*

# Algorithms in Numerical Linear Algebra (NLA)

For $Ax = b, Ax = \lambda(B)x, A = U\Sigma V^T$

1. **Classical** (dense) algorithms (LU, QR, Golub-Kahan)
   - $(+)$ Incredibly reliable, backward stable
   - $(-)$ Cubic complexity $O(n^3)$
2. **Iterative** (e.g. Krylov) algorithms
   - $(+)$ Fast convergence for 'good' matrices: clustered eigenvalues or (GMRES) or well-conditioned (LSQR)
   - $(-)$ If not, need preconditioner
3. **Randomized** algorithms
   - $(+)$ Next slide(s)
   - $(-)$ Lack of reproducibility, might lose nice properties, e.g. structure

# What can randomization do for you?

1. Sketch and **solve/precondition**
   - least-squares [Rokhlin-Tygert (08)], [Drineas-Mahoney-Muthukrishnan-Sarlós (10)], [Avron-Maymounkov-Toledo (10)], [Meng-Saunders-Mahoney 14]
2. **Near**-optimal solution with lightning speed
   - e.g. SVD [Halko-Martinsson-Tropp (11)], [Woodruff (14)]
3. **Sample** to approximate
   - Monte Carlo style; often comes with error estimates
   - e.g. matrix multiplication [Drineas-Kannan-Mahoney (06)], trace estimation [Avron-Toledo (11)], [Musco-Musco-Woodruff (20)]
4. Avoid pathological situations by perturbation/blocking
   - e.g. eigenvalues [Banks-Vargas-Kulkarni-Srivastava (19)], block Lanczos [Musco-Musco 15], [Tropp 18]

# What can randomization do for you?

1. Sketch and **solve/precondition** **Part I: linear/eigen solver**
   - least-squares [Rokhlin-Tygert (08)], [Drineas-Mahoney-Muthukrishnan-Sarlós (10)], [Avron-Maymounkov-Toledo (10)], [Meng-Saunders-Mahoney 14]
2. **Near**-optimal solution with lightning speed **Part II: low-rank SVD**
   - e.g. SVD [Halko-Martinsson-Tropp (11)], [Woodruff (14)]
3. **Sample** to approximate **(Part III: rank estimation)**
   - Monte Carlo style; often comes with error estimates
   - e.g. matrix multiplication [Drineas-Kannan-Mahoney (06)], trace estimation [Avron-Toledo (11)], [Musco-Musco-Woodruff (20)]
4. Avoid pathological situations by perturbation/blocking
   - e.g. eigenvalues [Banks-Vargas-Kulkarni-Srivastava (19)], block Lanczos [Musco-Musco 15], [Tropp 18]

# Sketching: Key idea in randomized linear algebra

Roughly: to solve a problem w.r.t. $\boxed{A}$, form random matrix $Y$

and work with $\boxed{Y^T A}$ (or sometimes $\boxed{Y^T A X}$)

Key insight: the sketch inherits $A$'s low-dimensional structure if present

Success stories in

- **Low-rank approximation** [Halko-Martinsson-Tropp 11, Woodruff 14, N. 20 etc]
- **Least-squares** [Rokhlin-Tygert 09, Avron-Maymounkov-Toledo 10]
- **Linear sytems and eigenvalue problems** [N.-Tropp 21]
- Rank estimation [Meier-N. 21]
- and many others

# Sketching for least-squares problems

For $A$: $n \times k$, $n \gg k$

$$\min_x \left\| \boxed{A}\,\boxed{x} - \boxed{b} \right\|_2 \quad \Rightarrow \quad \min_{\hat{x}} \left\| \boxed{SA}\,\boxed{\hat{x}} - \boxed{Sb} \right\|_2$$

With "reasonable/random" **sketch** $S \in \mathbb{C}^{s \times n}$ ($s > k$, say $s = 2k$),

$$(1-\epsilon)\|Av - b\|_2 \leq \|S(Av - b)\|_2 \leq (1+\epsilon)\|Av - b\|_2,$$

for some $\epsilon$ (not small, e.g. $\epsilon = \frac{1}{2}$) "subspace embedding". Hence the sketched solution $\hat{x}$ satisfies

$$\|A\hat{x} - b\|_2 \leq \frac{1+\epsilon}{1-\epsilon}\|Ax - b\|_2.$$

- if $\|Ax - b\|_2$ is small, $\hat{x}$ is a great solution!
- $SA$ in $O(nk \log n)$ cost: SRFT, HRFT [Cartis-Fiala-Shao 21], sparse sketch [Sarlos 06, Clarkson-Woodruff 17]

# Explaining why sketching works via M-P

Marchenko-Pastur: 'Rectangular random matrices are well-conditioned'

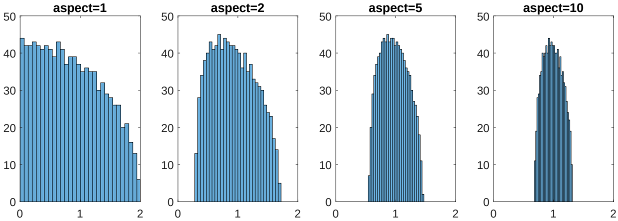


density $\sim \frac{1}{x}\sqrt{(\sqrt{m}+\sqrt{n})-x)(x-(\sqrt{m}-\sqrt{n})}$, support $[\sqrt{m}-\sqrt{n}, \sqrt{m}+\sqrt{n}]$

Claim: $\|Av-b\|_2 \approx \|S(Av-b)\|_2$ for all $v$ ($\approx$: 'same up to $O(1)$ factor')

- Let $[A,b]=QR$. $S[A,b]=(SQ)R$. Can write $\|Av-b\|_2=\|Qw\|_2$ and $\|S(Av-b)\|_2=\|(SQ)w\|_2$.
- Now $SQ$ is rectangular+random$\Rightarrow$ $\sigma_i(SQ)\approx 1$ by M-P.
- Hence $\|(SQ)w\|_2 \approx \|Qw\|_2$ for all $w$.

Related to J-L Lemma, RIP, oblivious subspace embedding etc

# GMRES for $Ax = b$ [Saad-Schulz 86]

Minimize residual in Krylov subspace $\mathcal{K}_d(A, b) := \text{span}(b, Ab, \ldots, A^{d-1}b)$

$$x_d = \text{argmin}_{x \in \mathcal{K}_d(A,b)} \|Ax - b\|_2$$

i.e., find solution of form $x_d = P_{d-1}(A)b$, $P_{d-1}$ polynomial deg $\leq d - 1$

# GMRES for $Ax = b$ [Saad-Schulz 86]

Minimize residual in Krylov subspace $\mathcal{K}_d(A,b) := \text{span}(b, Ab, \ldots, A^{d-1}b)$

$$x_d = \text{argmin}_{x \in \mathcal{K}_d(A,b)} \|Ax - b\|_2$$

i.e., find solution of form $x_d = P_{d-1}(A)b$, $P_{d-1}$ polynomial deg $\leq d-1$

Arnoldi process: finds orthonormal basis $Q_d$ for $\mathcal{K}_d(A,b)$

and write $x_d = Q_d y$, solve

$$\begin{aligned}\min_y \|AQ_d y - b\|_2 &= \min_y \|Q_{d+1}\tilde{H}_d y - b\|_2 \\ &= \min_y \left\| \begin{bmatrix} \tilde{H}_d \\ 0 \end{bmatrix} y - \|b\|_2 e_1 \right\|_2, \quad e_1 = [1, 0, \ldots, 0]^T\end{aligned}$$

- Reduces to Hessenberg least-squares, $O(d^2)$ work
- Overall, $d$ $A$-mult. $+\, O(nd^2)$ Arnoldi orthogonalization $+\, O(d^2)$ Hessenberg solve.
- Orthogonalization $O(nd^2)$ expensive→restarted GMRES

# Does sketching help in GMRES?

- Apparently NOT, as $\tilde{H}_d$ is almost square+full-rank in

$$\min_y \left\| \boxed{AQ_d}\,\boxed{y} - \boxed{b} \right\|_2 = \min_y \left\| \boxed{\tilde{H}_d}\,\boxed{y} - \boxed{e_1} \right\|_2$$

# Does sketching help in GMRES?

- Apparently NOT, as $\tilde{H}_d$ is almost square+full-rank in

$$\min_y \left\| AQ_d \, y - b \right\|_2 = \min_y \left\| \tilde{H}_d \, y - e_1 \right\|_2$$

- But $Q_d$ **need not be orthonormal!** Instead, we can find basis $B_d$ s.t. span$(B_d)$ = span$(Q_d)$ and solve

$$\min_y \left\| AB_d \, y - b \right\|_2$$

  - $x_d = B_d y$ is still the GMRES solution

# Does sketching help in GMRES?

- Apparently NOT, as $\tilde{H}_d$ is almost square+full-rank in

$$\min_y \left\| AQ_d \, y - b \right\|_2 = \min_y \left\| \tilde{H}_d \, y - e_1 \right\|_2$$

- But $Q_d$ **need not be orthonormal!** Instead, we can find basis $B_d$ s.t. span$(B_d) =$ span$(Q_d)$ and solve

$$\min_y \left\| AB_d \, y - b \right\|_2 \Rightarrow \min_{\hat{y}} \left\| SAB_d \, \hat{y} - Sb \right\|_2$$

  - $x_d = B_d y$ is still the GMRES solution
  - $AB_d$ is $n \times d$, ripe for sketching! Great if GMRES residual small

Does this buy us anything?

# Non-orthogonal linear algebra

$$\min_{y} \left\| AB_d\, y - b \right\|_2 \quad \Rightarrow \quad \min_{\hat{y}} \left\| SAB_d\, \hat{y} - Sb \right\|_2$$

- In GMRES, $B_d$ orthonormal; $O(nd^2)$ cost
- Not necessary! Works fine as long as $\kappa_2(B_d) < u^{-1} \approx 10^{16}$
- This offers enormous flexibility in choice of basis $B_d$, eliminating need for Arnoldi orthogonalization
  - Traditionally, want $\kappa_2(B_d) = 1 + O(\epsilon)$
  - Careful analysis (in Lanczos) shows $\kappa_2(B_d) = 1 + O(\sqrt{\epsilon})$ suffices
  - Sometimes, $\kappa_2(B_d) = O(1)$ is good enough
  - Here we're a lot more flexible; $B_d$ full rank $\kappa_2(B_d) \lesssim 10^{15}$
- Forming $[b, Ab, A^2b, \ldots]$ is still bad idea—explore alternatives

# Truncated orthogonalization

$b_1 = b/\|b\|_2$, and for $j = 2, 3, \ldots$, iteratively form

$$b_j = w_j/\|w_j\|_2 \quad \text{where} \quad w_j = (I - b_{j-1}b_{j-1}^* - \cdots - b_{j-k}b_{j-k}^*)(Ab_{j-1})$$

for a **fixed** $k$ (e.g. $k = 2, 4$)

- Orthogonalize only against last $k$ vectors
- When $A$ symmetric, reduces to Lanczos with $k = 2$
- Orthogonalization cost $O(nd)$ rather than $O(nd^2)$ after $d$ GMRES steps
- Often (not always) works quite well in practice wrt $\kappa_2(B_d)$ growth

# Krylov basis via Chebyshev recurrence

$$b_2 = Ab_1; \qquad b_i = 2Ab_{i-1} - b_{i-2} \quad \text{for } i = 3, \ldots, d.$$

- $b_i = T_{i-1}(A)b_1$, $T_i$ : Chebyshev polynomial
- We obtain $AB$ as a by-product.
- Shift+scale needed to adapt to the spectrum of $A$.
- No orthogonalization necessary—excellent efficiency (when it works)
- Block version trivially possible

# GMRES vs. sGMRES for $Ax = b$

$A$: $n \times n$, $d$: Krylov dimension, $k$: truncated orthogonalization para.
Typically $k \ll d \ll n$.

| | Matrix access | Form basis | Sketch | LS solve | Form soln. |
|---|---|---|---|---|---|
| Std. GMRES | $dT_{\mathrm{matvec}}$ | $nd^2$ | — | $d^2$ | $nd$ |
| sGMRES-$k$ | $dT_{\mathrm{matvec}}$ | $ndk$ | $nd \log d$ | $d^3$ | $nd$ |
| sGMRES-Cheb | $dT_{\mathrm{matvec}}$ | $nd$ | $nd \log d$ | $d^3$ | $nd$ |

# Experiments with sGMRES: FEM matrix

IFISS test problem (convection–diffusion)

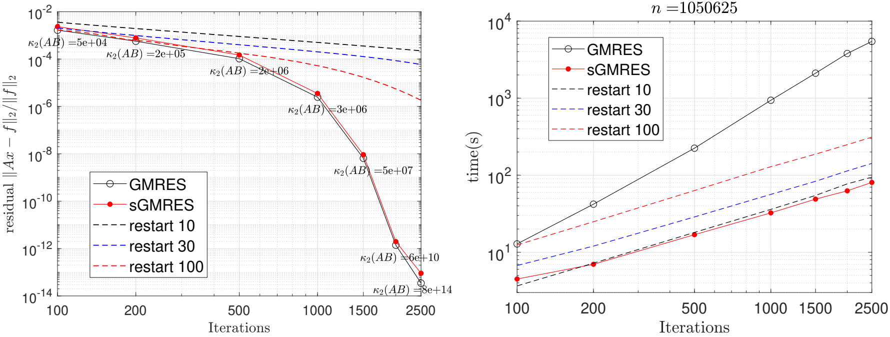


- sGMRES ($k = 4$-truncated) achieves 70x speedup over GMRES, comparable to GMRES-10 (restarted every 10 itereations)
- Accuracy of sGMRES is nearly identical to GMRES

# Experiments with sGMRES, PSD case

Discretized Laplacian

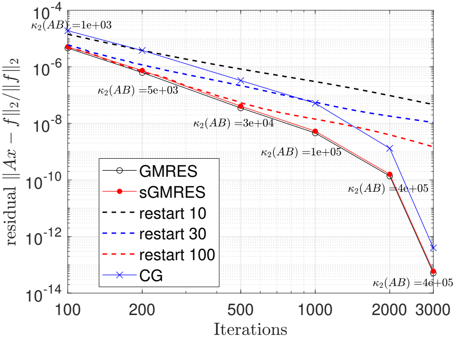


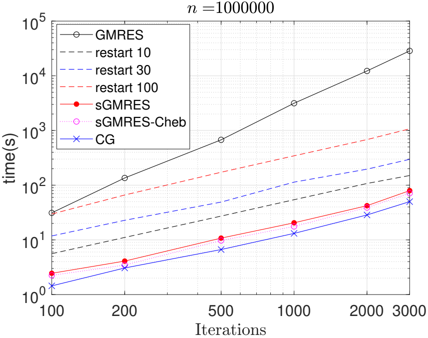


- Almost CG speed when $A \succ 0$!
- Of course, sGMRES does not require $A \succ 0$

"CG speed+GMRES flexibility"

# Cautionary note: ill-conditioning

Sometimes $\kappa_2(B)$ grows too fast: FS 680 1 from Matrix Market

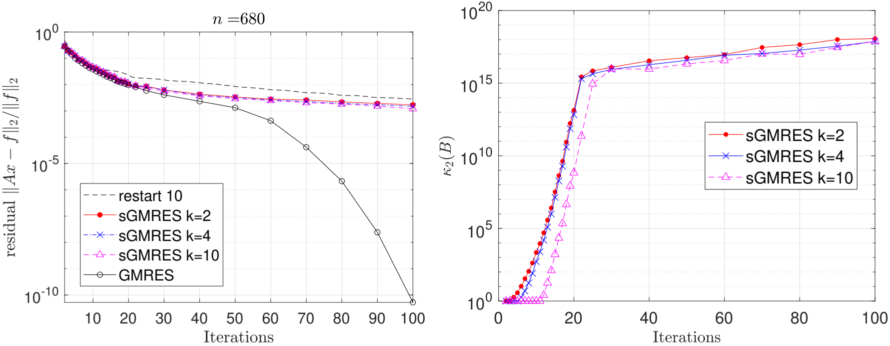


- Stagnation of sGMRES is purely a numerical issue
- Constructing full-rank basis $\kappa_2(B_d) < u^{-1}$ is crucial open problem
  - deflation
  - orthogonalize the sketch
  - restart once $\kappa_2(B_d) \gtrsim u^{-1}$ (rather than, say $d = 100$), etc.

# Eigenvalue problems: Rayleigh-Ritz

$$Ax = \lambda x$$

- If $A$ modest size $\lesssim 5000$, standard QR alg. is amazing
- For larger problems, **subspace methods**: find subspace $B \in \mathbb{R}^{n \times d}$ that approximately contains desired eigenvector(s), and solve.

# Eigenvalue problems: Rayleigh-Ritz

$$Ax = \lambda x$$

- If $A$ modest size $\lesssim 5000$, standard QR alg. is amazing
- For larger problems, **subspace methods**: find subspace $B \in \mathbb{R}^{n \times d}$ that approximately contains desired eigenvector(s), and solve.

Classical approach: **Rayleigh-Ritz** (RR); $B = QR$ (if $B$ not orthonormal), and solve small eigenproblem

$$Q^T A Q y = \lambda y.$$

$(\lambda, Qy)$ is approximate eigenpair (Ritz pair). Cost (at least) $O(nd^2)$

# Eigenvalue problems: Rayleigh-Ritz

$$Ax = \lambda x$$

- If $A$ modest size $\lesssim 5000$, standard QR alg. is amazing
- For larger problems, **subspace methods**: find subspace $B \in \mathbb{R}^{n\times d}$ that approximately contains desired eigenvector(s), and solve.

Classical approach: **Rayleigh-Ritz** (RR); $B = QR$ (if $B$ not orthonormal), and solve small eigenproblem

$$Q^T AQy = \lambda y.$$

$(\lambda, Qy)$ is approximate eigenpair (Ritz pair). Cost (at least) $O(nd^2)$
**Can we sketch Rayleigh-Ritz**? Doesn't look that way...
One can instead solve

$$B^T ABy = \lambda B^T By,$$

but still $O(nd^2)$ operations and stability issues.

# Alternative formulation of Rayleigh-Ritz

Key fact: RR is equivalent to solving $My = \lambda y$, where

$$\underset{M \in \mathbb{C}^{d\times d}}{\text{minimize}} \quad \left\| AB - B\,M \right\|_F ,$$

# Alternative formulation of Rayleigh-Ritz

Key fact: RR is equivalent to solving $My = \lambda y$, where

$$\underset{M \in \mathbb{C}^{d \times d}}{\text{minimize}} \left\| \boxed{AB} - \boxed{B}\,\boxed{M} \right\|_F ,$$

This can be sketched! **sRR** (sketched Rayleigh-Ritz)

$$\underset{\hat{M} \in \mathbb{C}^{d \times d}}{\text{minimize}} \left\| \boxed{SAB} - \boxed{SB}\,\boxed{\hat{M}} \right\|_F$$

Equivalent formulations:

- RR as *rectangular* eigenvalue problem $ABy \approx \lambda By$ [Ito-Murota 16]
  sketched version: $SABy \approx \lambda SBy$
- RR as Galerkin orthogonalization $B^T(ABy - \lambda By) = 0$; sketched version: $B^T S^T S(ABy - \lambda By) = 0$

# Possible bases $B_d$ for eigenvalue problems

- Krylov subspace
- Block Krylov subspace [Musco-Musco (15)], [Tropp (18)] etc
  - Chebyshev/Newton recurrence appealing
- LOBPCG ($B$ via previous block update and current block descent), Jacobi-Davidson ($B$ via linear systems), etc.
- Eigenspace for nearby eigenproblems

We'll show examples with Krylov (to compare with eigs)—but other choices can be more appealing

# Experiments: nonsymmetric eigs

$B$: Krylov (for comparison with eigs)

Nonsymmetric eigenproblem arising in trust-region subproblem (from [Rojas-Santos-Sorensen (08)])rightmost eigenpair desired

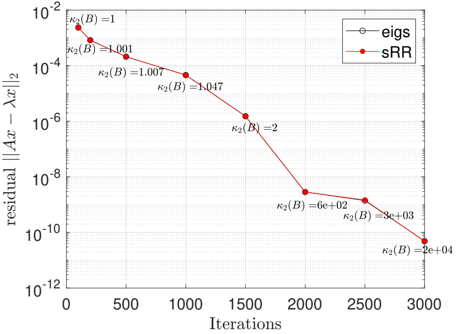


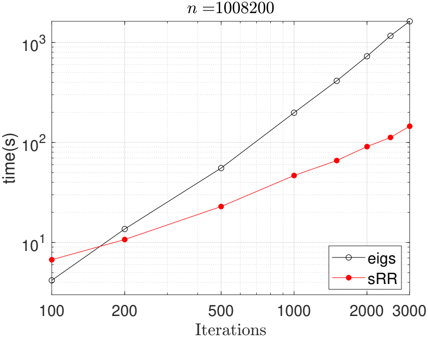

# Experiments: symmetric eigs via block Lanczos

Laplacian matrix, find (smallest) eigenpairs via block Krylov subspace with block size $b = 10$.

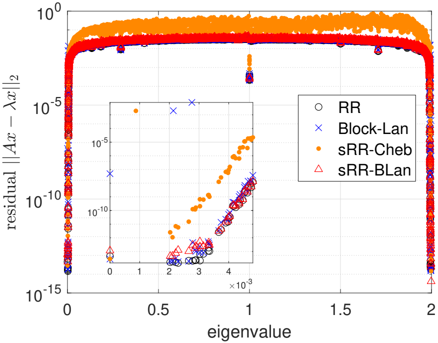


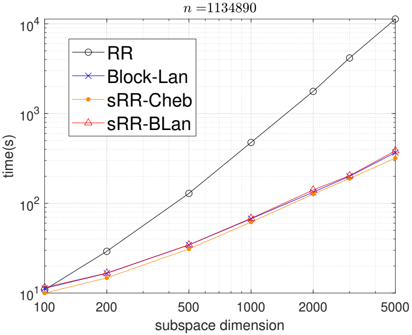


Block-Lan: block Lanczos (without full orthogonalization)

sRR-Cheb: build subspace via block Chebyshev recurrence

sRR-Blan: use block Lanczos subspace for sRR

sRR largely avoids "ghost eigenvalues"

# Summary of Part I

- Sketching is VERY useful
- sGMRES offers new opportunities+challenges
  - Goals of preconditioner
  - Reconsider restart strategies
  - Building full-rank basis $B_d$
- sRR can be applied for
  - $Ax = \lambda Bx$, SVD, rectangular eigenproblem

Opportunities&Problems

- Block (Krylov) methods appealing for eig/SVD on parallal computers
- Detecting/avoiding rank-deficiency of $B_d$
- New way of looking at preconditioning/restarting?
- Opportunities with non-orthogonal linear algebra?

# (Most) important result in Numerical Linear Algebra

Given $A \in \mathbb{R}^{m\times n}$ $(m \geq n)$, find low-rank (rank $r$) approximation

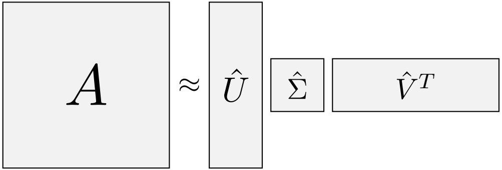


, $\quad \hat{\Sigma} \in \mathbb{R}^{r\times r}$

- Optimal solution $A_r = U_r \Sigma_r V_r^T$ via truncated SVD
  $U_r = U(:, 1:r), \Sigma_r = \Sigma(1:r, 1:r), V_r = V(:, 1:r)$, giving

$$\|A - A_r\| = \|\mathrm{diag}(\sigma_{r+1}, \ldots, \sigma_n)\|$$

  in any unitarily invariant norm [Horn-Johnson 1985]
- But that costs $O(mn^2)$; look for faster approximation
- Low-rank matrices everywhere [Beckermann-Townsend 17, Gillis 20 etc]

# Part II: Randomized low-rank approximation

[Halko-Martinsson-Tropp, SIREV 2011]

1. Form a random matrix $X \in \mathbb{R}^{n \times r}$.
2. Compute $AX$.
3. QR factorization $AX = QR$.
4. $A \approx QQ^T A =: \hat{A} = (QU_0)\Sigma_0 V_0^T$ is approximate SVD.

- $O(mnr)$ cost for dense $A$, can be reduced to $O(mn \log n + mr^2)$ via FFT and interp. decomp. (slightly worse accuracy)
- $mr^2$ dominant if $r > \sqrt{n}$ or e.g. $A$ sparse
- Near-optimal approximation guarantee: for any $\hat{r} < r$,

$$\mathbb{E}\|A - \hat{A}\|_F \leq \left(1 + \frac{r}{r - \hat{r} - 1}\right) \|A - A_{\hat{r}}\|_F$$

where $A_{\hat{r}}$ is the (optimal) rank $\hat{r}$-truncated SVD

# Approximants of form $AX(Y^TAX)^\dagger Y^TA$

**Generalized Nyström** ( GN) for general $A \in \mathbb{R}^{m\times n}$:

$$A \approx AX(Y^TAX)^\dagger Y^TA =$$

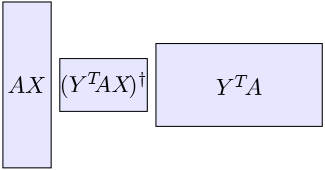


- $X \in \mathbb{R}^{n\times r}, Y \in \mathbb{R}^{m\times(r+\ell)}$, $\ell = cr$ (we choose $c = 0.5$)
  - e.g. **Gaussian** $X_{ij} \sim N(0,1)$
  - or **SRFT** $X = DFS$, $D$ : diag, $F$: FFT, $S$: subsampling (or hashing)
- Near-optimal cost, essentially $AX$ and $Y^TA$. Single-pass
- Near-optimal accuracy, comparable to HMT, Nyström

# Approximants of form $AX(Y^TAX)^\dagger Y^TA$

**stabilized Generalized Nyström** (SGN) for general $A \in \mathbb{R}^{m\times n}$:

$$A \approx AX(Y^TAX)^\dagger_\epsilon Y^TA =$$

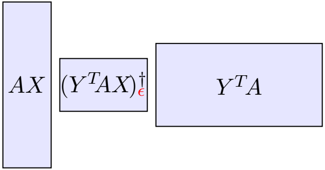


- $X \in \mathbb{R}^{n\times r}, Y \in \mathbb{R}^{m\times(r+\ell)}$, $\ell = cr$ (we choose $c = 0.5$)
  - e.g. **Gaussian** $X_{ij} \sim N(0,1)$
  - or **SRFT** $X = DFS$, $D$ : diag, $F$: FFT, $S$: subsampling (or hashing)
- Near-optimal cost, essentially $AX$ and $Y^TA$. Single-pass
- Near-optimal accuracy, comparable to HMT, Nyström
- Numerically stable with $\epsilon$-pseudoinverse $(U\Sigma V^T)^\dagger_\epsilon = V\Sigma^\dagger_\epsilon U^T$

# Approximants of form $AX(Y^TAX)^\dagger Y^TA$

**stabilized Generalized Nyström** (SGN) for general $A \in \mathbb{R}^{m\times n}$:

$$A \approx AX(Y^TAX)^\dagger_\epsilon Y^TA = \boxed{AX}\,\boxed{(Y^TAX)^\dagger_\epsilon}\,\boxed{Y^TA}$$

- $X \in \mathbb{R}^{n\times r}, Y \in \mathbb{R}^{m\times(r+\ell)}$, $\ell = cr$ (we choose $c = 0.5$)
  - e.g. **Gaussian** $X_{ij} \sim N(0,1)$
  - or **SRFT** $X = DFS$, $D$ : diag, $F$: FFT, $S$: subsampling (or hashing)
- Near-optimal cost, essentially $AX$ and $Y^TA$. Single-pass
- Near-optimal accuracy, comparable to HMT, Nyström
- Numerically stable with $\epsilon$-pseudoinverse $(U\Sigma V^T)^\dagger_\epsilon = V\Sigma^\dagger_\epsilon U^T$
- Key tool for convergence+stability analysis: **Marchenko-Pastur**

  "rectangular random matrices are well conditioned"

  Related to J-L Lemma, RIP, oblivious subspace embedding etc

# Approximants of form $AX(Y^TAX)^\dagger Y^TA$

(or $A(A^TA)^qX(Y^TA(A^TA)^qX)^\dagger Y^TA$)

$\Omega$: random matrix (e.g. Gaussian, SRFT)

| | $X, Y$ | $q$ | stable? | cost for dense $A$ |
|---|---|---|---|---|
| HMT 2011 | $X = \Omega, Y = AX$ | 0 | $(\surd)$ | $O(mnr)$ |
| Nyström ($A \succ 0$) | $Y = X = \Omega$ | 0 | $(\times)$ | $O(mn\log n + mr^2)$ |
| HMT+Nyström | $Y = X = Q, A\Omega = QR$ | 1 | $(\times)$ | $O(mnr)$ |
| Subspace iter | $X = \Omega, Y = \tilde{\Omega}$ | $> 1$ | $(\surd)$ | $O(mnrq)$ |
| TYUC19 | (4 sketch matrices) | 0 | $(\surd)$ | $O(mn\log n + mr^2)$ |
| TYUC17 | $X = \Omega, Y = \tilde{\Omega}$ | 0 | $(\surd)$ | $O(mn\log n + mr^2)$ |
| Clarkson-Woodruff09(C-W) | $X = \Omega, Y = \tilde{\Omega}$ | 0 | $(\times)$ | $O(mn\log n + r^3)$ |
| Demmel-Grigori-Rusciano19 | C-W+extra term | 0 | $(\times)$ | $O(mn\log n + mr^2)$ |
| **This work, GN** | $X = \Omega, Y = \tilde{\Omega}$ | 0 | $\surd$ | $O(mn\log n + r^3)$ |

$(\times)$: unstable examples exist (though often perform ok)

$(\surd)$: conjectured to be stable (no proof)

- GN Combines **stability** and **near-optimal complexity**
- explicit constants available: GN $10mn\log n + \frac{7}{3}r^3$ flops

# Experiments: dense matrix

Dense $50000 \times 50000$ matrix w/ geom. decaying $\sigma_i$

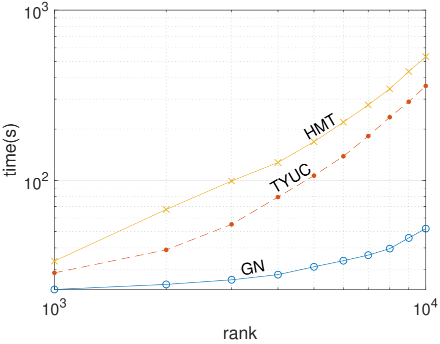


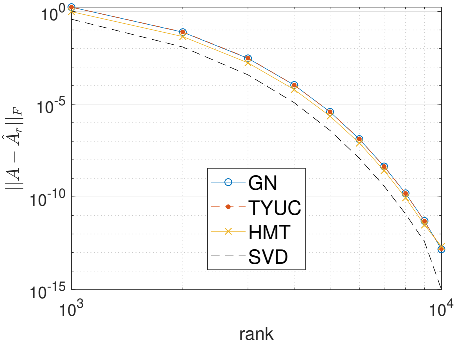


HMT: Halko-Martinsson-Tropp 11, TYUC: Tropp-Yurtsever-Udell-Cevher 17

- GN and TYUC have same accuracy (as they should)
- GN faster, up to $\approx 10$x

# Part II in a nutshell

```
n = 1000; % size
A = gallery('randsvd',n,1e100);
r = 200;  % rank

X = randn(n,r); Y = randn(n,1.5*r);
AX = A*X;
YA = Y'*A;
YAX = YA*X;
[Q,R] = qr(YAX,0);    % stable implementation of pseudoinverse
At = (AX/R)*(Q'*YA);

norm(At-A,'fro')/norm(A,'fro')
ans = 2.8138e-15
```

For details, please see my E-NLA talk, and arXiv 2009.11392
"Fast and stable randomized low-rank matrix approximation"
Also related: "Randomized low-rank approximation for symmetric indefinite matrices", arXiv 2212.01127 (with T. Park)

# Summary

Randomization can help you with:

1. **Sketch** and solve/precondition
   - **Part I: linear(/eigen) solver**
2. **Near**-optimal solution with lightning speed
   - **Part II: low-rank SVD**
3. **Sample** to approximate
   - **(Part III: rank estimation)** (if time permits..)
4. Avoid bad situations by perturbation/blocking

# Part I: Rank estimation

In most low-rank algorithms, the rank $r$ is required as input

- If $r$ too low: need to resample $A$ and recompute
- If $r$ too high: wasted computation

A fast rank estimator is thus highly desirable

**Definition**

$\mathrm{rank}_\epsilon(A)$: integer $i$ s.t. $\sigma_i(A) > \epsilon \geq \sigma_{i+1}(A)$.

This work: $O(mn \log n + r^3)$ algorithm for rank estimation

[with Maike Meier (Oxford), arXiv 2021]

- In many cases, extra cost is much lower (e.g. $O(r^2)$)
- Key idea: **Sample** the singular values via sketching, $Y^T AX$

# Goal of a rank estimator

It is usually not necessary (or even possible, with subcubic work) to find the exact $\epsilon$-rank.

We aim to find $\hat{r}$ s.t.

- $\sigma_{\hat{r}+1}(A) = O(\epsilon)$ (say, $\sigma_{\hat{r}+1}(A) < 10\epsilon$): $\hat{r}$ is not a severe underestimate, and
- $\sigma_{\hat{r}}(A) = \Omega(\epsilon)$ (say, $\sigma_{\hat{r}}(A) > 0.1\epsilon$): $\hat{r}$ is not a severe overestimate.

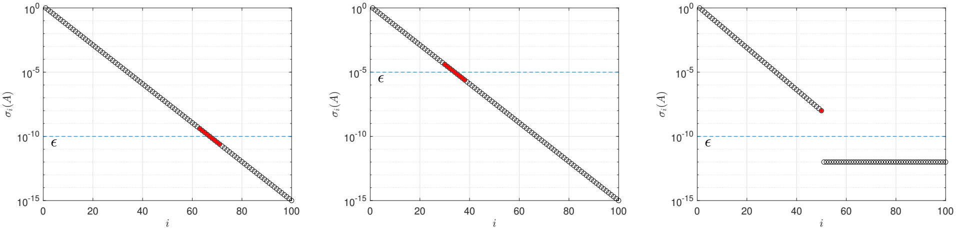

# Goal of a rank estimator

It is usually not necessary (or even possible, with subcubic work) to find the exact $\epsilon$-rank.

We aim to find $\hat{r}$ s.t.

- $\sigma_{\hat{r}+1}(A) = O(\epsilon)$ (say, $\sigma_{\hat{r}+1}(A) < 10\epsilon$): $\hat{r}$ is not a severe underestimate, and
- $\sigma_{\hat{r}}(A) = \Omega(\epsilon)$ (say, $\sigma_{\hat{r}}(A) > 0.1\epsilon$): $\hat{r}$ is not a severe overestimate.

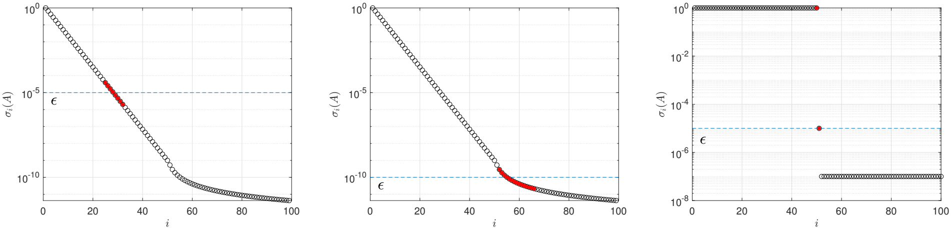


Consequently, it suffices to estimate $\sigma_i(A)$ to their order of magnitude

# Previous studies on rank estimation

- Based on full factorization (e.g. Duersch-Gu 2020, Martinsson-Quintana-Orti-Heavner 2019)
  - cubic $O(mn^2)$complexity
- Ubaru-Saad (2016): polynomial approximation and spectral density estimates using Krylov subspace methods
  - complexity difficult to predict
- Andoni-Nguyen (2013): theory that suggest rankest possible, no algorithm

Our algorithm: based on random sketches $AX$, $Y^TAX$

Key fact: $\sigma_i(AX)/\sigma_i(A) = O(1)$ for leading $i$, and $\sigma_i(Y^TAX)/\sigma_i(AX) = O(1)$

- Study of $\sigma_i(AX)$ is covariance estimate
  - Usually, at least $n$ samples required
  - But **leading** sing vals good with many fewer samples

# Main idea: random embedding preserves $O(\sigma_i)$

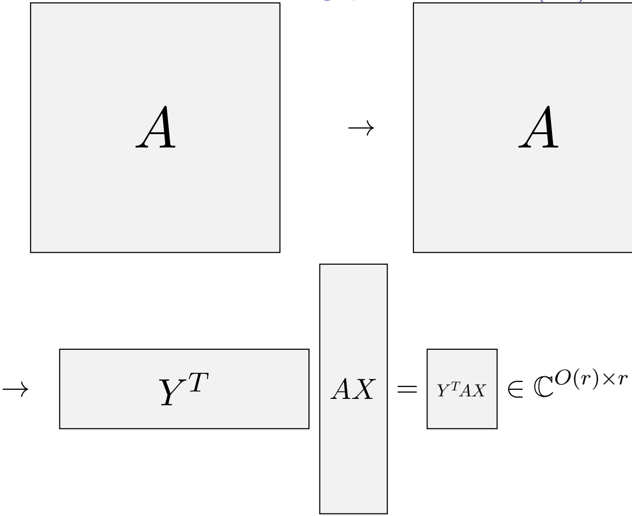


$X, Y$: Gaussian (or SRFT), scaled s.t. $\sigma_i(Q^TX), \sigma_i(YQ) \in [1-\delta, 1+\delta]$.

Key fact: $\frac{\sigma_i(A)}{\sigma_i(Y^TAX)} = O(1)$ for $i = 1, 2, \ldots, r$

# $\sigma_i(AX)/\sigma_i(A) = O(1)$ for leading $i$

Let $G \in \mathbb{C}^{n\times r}$ and

$$AG = U_1\Sigma_1(V_1^*G) + U_2\Sigma_2(V_2^*G) = U_1\Sigma_1G_1 + U_2\Sigma_2G_2,$$

**Lemma**

*For $i = 1, \dots, r$,*

$$\sigma_{\min}(\hat{G}_{\{i\}}) \le \frac{\sigma_i(AG)}{\sigma_i(A)} \le \sqrt{\sigma_{\max}(\tilde{G}_{\{r-i+1\}})^2 + \left(\frac{\sigma_{r+1}(A)\sigma_{\max}(G_2)}{\sigma_i(A)}\right)^2}$$

*$\hat{G}_{\{i\}} \in \mathbb{C}^{i\times r}$: first $i$ rows of $G_1$, and $\tilde{G}_{\{r-i+1\}}$ last $r-i+1$ rows of $G_1$. If $G$ is standard Gaussian, $\hat{G}_{\{i\}}$, $\tilde{G}_{\{r-i+1\}}$, and $G_2$ are independent standard Gaussian.*

PROOF: Courant-Fisher minimax characterization.

# $\sigma_i(AX)/\sigma_i(A) = O(1)$ cont'd

$$\sigma_{\min}(\hat{G}_{\{i\}}) \le \frac{\sigma_i(AG)}{\sigma_i(A)} \le \sqrt{\sigma_{\max}(\tilde{G}_{\{r-i+1\}})^2 + \left(\frac{\sigma_{r+1}(A)\sigma_{\max}(G_2)}{\sigma_i(A)}\right)^2}$$

When $X$ **scaled Gaussian** (embedding)

**Theorem**

*Let* $X \in \mathbb{R}^{n\times r}$ *with* $X_{ij} \sim N(0, 1/r)$. *Then for* $i = 1, \dots, r$

$$1 - \sqrt{\frac{i}{r}} \le \mathbb{E}\frac{\sigma_i(AX)}{\sigma_i(A)} \le 1 + \sqrt{\frac{r-i+1}{r}} + \frac{\sigma_{r+1}}{\sigma_i}\left(1 + \sqrt{\frac{n-r}{r}}\right).$$

*Failure probability decays squared-exponentially*

Proof: Marchenko-Pastur ("rectangular random matrices are well-conditioned")

- Interpretation: $\frac{\sigma_i(AX)}{\sigma_i(A)} \approx 1$, esp. for small $r$

# $\sigma_i(AX)/\sigma_i(A) = O(1)$ cont'd

$$\sigma_{\min}(\hat{G}_{\{i\}}) \le \frac{\sigma_i(AG)}{\sigma_i(A)} \le \sqrt{\sigma_{\max}(\tilde{G}_{\{r-i+1\}})^2 + \left(\frac{\sigma_{r+1}(A)\sigma_{\max}(G_2)}{\sigma_i(A)}\right)^2}$$

When $X$ general embedding

**Theorem**

*Let $\tilde{V}_1$ be $A$'s top right singvecs, and suppose $\sigma_i(V_1^T X) \in [1-\epsilon, 1+\epsilon]$ for some $\epsilon < 1$. Then, for $i = 1, \dots, \tilde{r}$*

$$1 - \epsilon \le \frac{\sigma_i(AX)}{\sigma_i(A)} \le \sqrt{(1+\epsilon)^2 + \left(\frac{\sigma_{\tilde{r}+1}(A)\|X\|_2}{\sigma_i(A)}\right)^2}.$$

$\epsilon$-subspace embedding, (e.g. SRFT (subsampled random Fourier transform), i.e. $X = DFS$, $D$ : diag, $F$: FFT, $S$: subsampling), also effective choices for $X$

# Experiments $\sigma_i(AX)/\sigma_i(A) = O(1)$

$A \in \mathbb{R}^{1000\times 1000}$

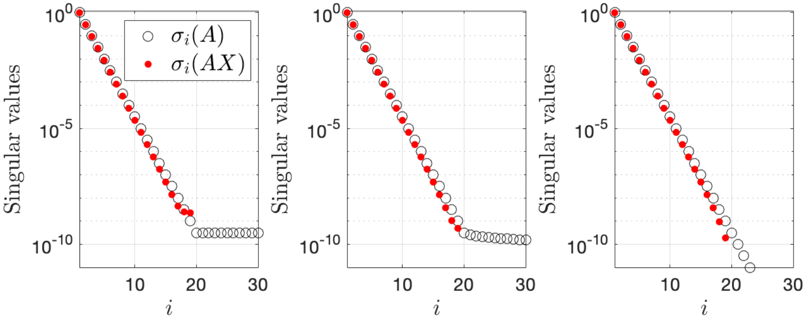


- Leading singvals estimated reliably (when they decay)
- Tail effect nonnegligible (esp. for last $i \approx r$)
- Hence trust only leading (say $90\%$) samples

# 2nd step: $\sigma_i(Y^T AX)/\sigma_i(AX) = O(1)$

**Corollary (Combines Boutsidis-Gittens (13) and Tropp (11))**

*Let* $AX \in \mathbb{R}^{m\times r_1}$*, with* $m \geq r_1$*, and let* $Y \in \mathbb{R}^{n\times r_2}$ *be an SRFT matrix. Let* $0 < \epsilon < 1/3$ *and* $0 < \delta < 1$*. If*

$$r_2 \geq 6\eta\epsilon^{-2}\left[\sqrt{r_1} + \sqrt{8\log(m/\delta)}\right]^2 \log(r_1/\delta),$$

*then with failure probability at most* $3\delta$

$$\sqrt{1-\epsilon} \leq \frac{\sigma_i(Y^T AX)}{\sigma_i(AX)} \leq \sqrt{1+\epsilon},$$

*for each* $i = 1, \ldots, r_1$.

# $\sigma_i(Y^TAX)/\sigma_i(AX) = O(1)$

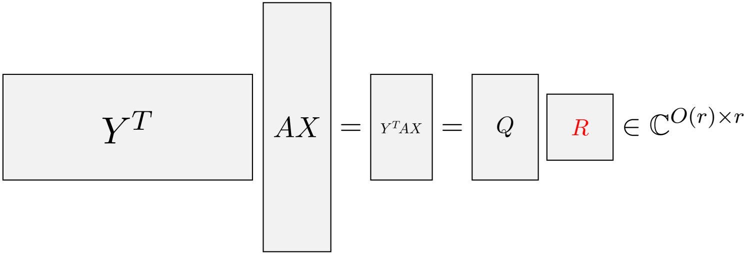


- Approximate orthogonalization: ideas from Blendenpik etc [Avron-Maymounkov-Toledo 10]
- In generalized Nyström, $Y^TAX = QR$ already computed + rank-revealing QR $\Rightarrow \sigma_i(Y^TAX) \approx \text{diag}(R)$; only $O(r)$ extra cost

# Experiments: $\sigma_i(Y^TAX)/\sigma_i(AX) = O(1)$

$AX \in \mathbb{R}^{10^5\times 2000}$

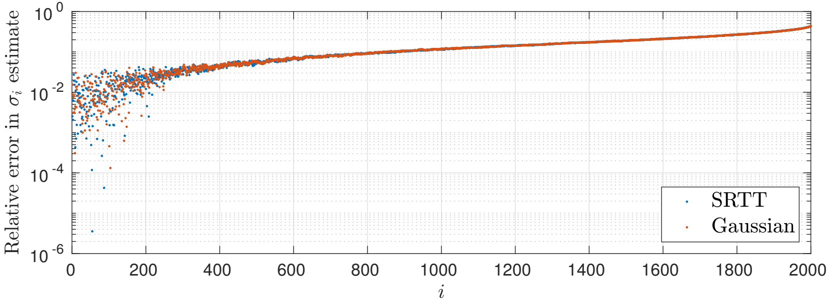


- $|\frac{\sigma_i(Y^TAX)}{\sigma_i(AX)} - 1|$ small esp. for leading singvals
- Reasonable estimates even for $i \approx r$

# The rank estimation algorithm

**Algorithm** Given $A \in \mathbb{C}^{m \times n}$, tolerance $\epsilon$ and an upper bound for rank $r_1$, compute approximate $\epsilon$-rank.

1: Set $\tilde{r}_1 = \text{round}(1.1 r_1)$ to oversample by 10%.
2: Draw $n \times \tilde{r}_1$ random embedding matrix $X$.
3: Form the $m \times \tilde{r}_1$ matrix $AX$.

**2. Approximate orthogonalization:**

4: Set $r_2 = 1.5\tilde{r}_1$, draw an $r_2 \times m$ SRFT embedding matrix $Y$.
5: Form the $r_2 \times \tilde{r}_1$ matrix $Y^T AX$.

**3. Singular value estimates:**

6: Compute the first $r_1$ singular values of $Y^T AX$.
7: Output smallest $\hat{r}$ s.t. $\sigma_{\hat{r}+1}(Y^T AX) \le \epsilon$.

Complexity: $O(mn \log n + r^3)$

# Experiments: rank estimation

SP/FP: slow/fast polynomial decay in $\sigma_i(A)$, SE/FE: slow/fast exponential decay

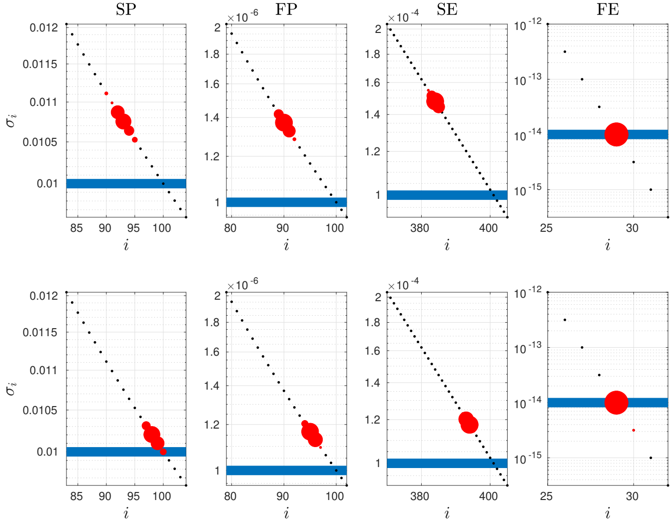


Out of 100 runs; dot area reflects frequency

# Experiments: gaps in singular values

$A_{G,IC}$: incoherent singvecs, $A_{G,C}$: coherent singvecs $(V = I)$

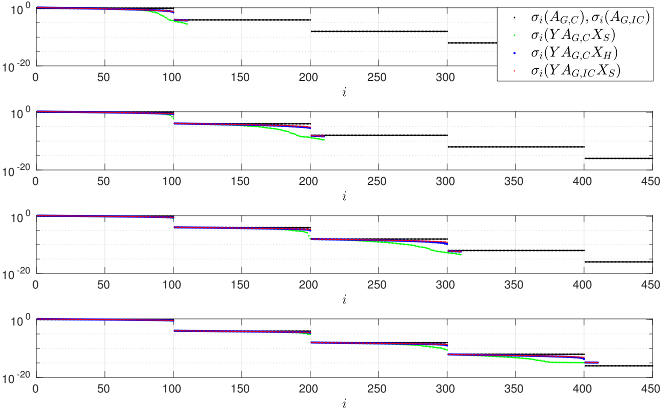


For coherent problems, *Hashed* (not subsampled) RFT helpful [Cartis-Fiala-Shao 21]

For details, please see preprint Meier-N. "Fast randomized numerical rank estimation" arXiv 2021.