Randomized Algorithms for Very Large-Scale Linear Algebra

Gunnar Martinsson The University of Colorado at Boulder

Collaborators: Edo Liberty, Vladimir Rokhlin (congrats!), Yoel Shkolnisky, Arthur Szlam, Joel Tropp, Mark Tygert, Franco Woolfe, ...

Student: Nathan Halko (now at Spot Influence, LLC).

Paper: N. Halko, P.G. Martinsson, J. Tropp. *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions.* SIAM Review, **53**(2), 2011.

Slides: Posted under the "Talks" tab on my webpage (google "Gunnar Martinsson").

Objective:

Given an $n \times n$ matrix **A**, for a very large *n* (say $n \sim 10^6$), we seek to compute a rank-*k* approximation, with $k \ll n$ (say $k \sim 10^2$ or 10^3),

$$\mathbf{A} \approx \mathbf{E} \mathbf{F}^* = \sum_{j=1}^k \mathbf{e}_j \mathbf{f}_j^*.$$

 $n \times n$ $n \times k$ $k \times n$

Solving this problem leads to algorithms for computing:

- Eigenvectors corresponding to leading eigenvalues. (Require $\mathbf{e}_j = \lambda_j \mathbf{f}_j$, and $\{\mathbf{f}_j\}_{j=1}^k$ to be orthonormal.)
- Singular Value Decomposition (SVD) / Principal Component Analysis (PCA).
 (Require {e_j}^k_{j=1} and {f_j}^k_{j=1} to be orthogonal sets.)
- Spanning columns or rows.

(Require $\{e_j\}_{j=1}^k$ to be columns of **A**, or require $\{f_j^*\}_{j=1}^k$ to be rows of **A**.)

• etc

The problem being addressed is ubiquitous in applications.

Note: The only error that will be controlled is the "backwards error" $||\mathbf{A} - \mathbf{EF}|| \le \text{TOL}$.

Applications:

- Fast algorithms for elliptic PDEs: more efficient Fast Multipole Methods, fast *direct* solvers, construction of special quadratures for corners and edges, etc.
- Statistical analysis via Principal Component Analysis.
- Data mining (machine learning, analysis of network matrices, etc).
- Diffusion geometry; a technique for constructing parameterizations on large collections of data points organized (modulo noise) along non-linear low-dimensional manifolds.
 Requires the computations of eigenvectors of *graph Laplace operators*.
- Nearest neighbor search for large clouds of points in high dimensional space.
- General pre-conditioners.
- Etc.

Review of existing methods I

For a dense $n \times n$ matrix that fits in RAM, excellent algorithms are already part of LAPACK (and Matlab, *etc*).

- Double precision accuracy.
- Very stable.
- $O(n^3)$ asymptotic complexity. Reasonably small constants.
- Require extensive random access to the matrix.

When the target rank k is much smaller than n, there also exist $O(n^2 k)$ methods with similar characteristics (the well-known Golub-Businger method, RRQR by Gu and Eisentstat, *etc*).

For small^a matrices the state-of-the-art is very satisfactory. For kicks, we will improve on it anyway, but this is not the main point.

^aOn a standard 2012 personal computer, "small" means $n \le 10\,000$ or so.

Review of existing methods II

If the matrix is large, but can rapidly be applied to a vector (if it is sparse, or sparse in Fourier space, or amenable to the FMM, etc.), so called *Krylov subspace methods* often yield excellent accuracy and speed.

The idea is to pick a starting vector ω (often a random vector), "restrict" the matrix **A** to the *k*-dimensionsal "Krylov subspace"

$$\mathsf{Span}(oldsymbol{\omega},\, \mathbf{A}\,oldsymbol{\omega},\, \mathbf{A}^2\,oldsymbol{\omega},\, \dots,\, \mathbf{A}^{k-1}\,oldsymbol{\omega})$$

and compute approximate eigenvectors of the resulting matrix. Advantages:

- Very simple access to A.
- Extremely high accuracy possible.

Drawbacks:

- In standard implementations, the matrix is revisited O(k) times if a rank-k approximation is sought. (Blocked versions exist, but the convergence analysis is less developed.)
- There are numerical stability issues. These are well-studied and can be overcome, but they make the algorithms less portable (between applications, hardware platforms, etc.).

"New" challenges in algorithmic design:

The existing state-of-the-art methods of numerical linear algebra that we have very briefly outlined were designed for an environment where the matrix fits in RAM and the key to performance was to minimize the number of floating point operations required.

Currently, *communication* is becoming the real bottleneck:

- While clock speed is hardly improving at all anymore, the cost of a flop keeps going down rapidly. (Multi-core processors, GPUs, cloud computing, etc.)
- The cost of slow storage (hard drives, flash memory, etc.) is also going down rapidly.
- Communication costs are decreasing, but *not* rapidly.
 - Moving data from a hard-drive.
 - Moving data between nodes of a parallel machine. (Or cloud computer ...)
 - The amount of fast cache memory close to a processor is not improving much.
 (In fact, it could be said to be *shrinking* GPUs, multi-core, etc.)
- "Deluge of data". Driven by ever cheaper storage and acquisition techniques. Web search, data mining in archives of documents or photos, hyper-spectral imagery, social networks, gene arrays, proteomics data, sensor networks, financial transactions, ...

The more powerful computing machinery becomes, the more important efficient algorithm design becomes.

- Linear scaling (w.r.t. problem size, processors, etc.).
- Minimal data movement.

That *randomization* can be used to overcome some of the communication bottlenecks in matrix computations has been pointed out by several authors:

C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala (2000)

A. Frieze, R. Kannan, and S. Vempala (1999, 2004)

D. Achlioptas and F. McSherry (2001)

P. Drineas, R. Kannan, M. W. Mahoney, and S. Muthukrishnan (2006a, 2006b, 2006c, 2006d, *etc*)

S. Har-Peled (2006)

- A. Deshpande and S. Vempala (2006)
- S. Friedland, M. Kaveh, A. Niknejad, and H. Zare (2006)
- T. Sarlós (2006a, 2006b, 2006c)
- K. Clarkson, D. Woodruff (2009)

Literature survey: Halko, Martinsson, Tropp (2011).

Review of existing methods III

Examples of how randomization could be used:

• Random column/row selection

Draw at random some columns and suppose that they span the entire column space. If rows are drawn as well, then spectral properties can be estimated. Crude sampling leads to less than O(mn) complexity, but is very dangerous.

• Sparsification

Zero out the vast majority of the entries of the matrix. Keep a random subset of entries, and boost their magnitude to preserve "something."

Quantization and sparsification

Restrict the entries of the matrix to a small set of values (-1/0/1 for instance).

The methods outlined can be as fast as you like, but must necessarily have very weak performance guarantees. They can work well for certain classes of matrices for which additional information is available (basically, matrices that are in some sense "over-sampled").

Approach advocated here:

A randomized algorithm for computing a rank-k approximation to an $m \times n$ matrix. It is engineered from the ground up to:

- Minimize communication.
- Handle streaming data, or data stored "out-of-core."
- Easily adapt to a broad range of distributed computing architectures.

Computational profile:

- At least O(mn) complexity (in fact O(mnk) or $O(mn\log(k))$).
- The accuracy ε is a user-set number. (If the application permits, it could be $\varepsilon = 10^{-12}$ or less.)
- Since the method is randomized, it has a *failure probability* η.
 η is a user specified number.

The cost of the method grows as $\eta \to 0$, but setting $\eta = 10^{-10}$ is cheap. For all practical purposes, the methods succeed with probability 1. **Goal:** Given an $m \times n$ matrix **A**, compute an approximate rank-k SVD **A** \approx **U** Σ **V**^{*}.

Goal: Given an $m \times n$ matrix **A**, compute an approximate rank-k SVD $\mathbf{A} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$.

Algorithm:

- Find an m × k orthonormal matrix Q such that A ≈ Q Q* A.
 (*I.e.*, the columns of Q form an ON-basis for the range of A.)
 3.
- 4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
- 5. Compute the SVD of the small matrix **B**: $\mathbf{B} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \mathbf{V}^*$.
- 6. Form the matrix $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.

Goal: Given an $m \times n$ matrix **A**, compute an approximate rank-k SVD $\mathbf{A} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$.

Algorithm:

1.

2.

3.

Find an $m \times k$ orthonormal matrix **Q** such that $\mathbf{A} \approx \mathbf{Q} \mathbf{Q}^* \mathbf{A}$. (I.e., the columns of **Q** form an ON-basis for the range of **A**.)

- 4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
- 5. Compute the SVD of the small matrix **B**: $\mathbf{B} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \mathbf{V}^*$.
- 6. Form the matrix $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.

Note: Steps 4 – 6 are exact; the error in the method is all in **Q**:

$$||\mathbf{A} - \bigcup_{=\mathbf{Q}\hat{\mathbf{U}}} \mathbf{\Sigma}\mathbf{V}^*|| = ||\mathbf{A} - \mathbf{Q}\underbrace{\hat{\mathbf{U}}\mathbf{\Sigma}\mathbf{V}^*}_{=\mathbf{B}}|| = ||\mathbf{A} - \mathbf{Q}\underbrace{\mathbf{B}}_{\mathbf{Q}^*\mathbf{A}}|| = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||.$$

Note: The classical *Golub-Businger algorithm* follows this pattern. It finds **Q** in Step 3 via direct orthogonalization of the columns of **A** via, e.g., Gram-Schmidt.

Range finding problem: Given an $m \times n$ matrix **A** and an integer $k < \min(m, n)$, find an orthonormal $m \times k$ matrix **Q** such that $\mathbf{A} \approx \mathbf{Q}\mathbf{Q}^*\mathbf{A}$.

Solving the primitive problem via randomized sampling — intuition:

- 1. Draw random vectors $\omega_1, \omega_2, \ldots, \omega_k \in \mathbb{R}^n$. (We will discuss the choice of distribution later — think Gaussian for now.)
- 2. Form "sample" vectors $\mathbf{y}_1 = \mathbf{A} \, \boldsymbol{\omega}_1, \, \mathbf{y}_2 = \mathbf{A} \, \boldsymbol{\omega}_2, \, \dots, \, \mathbf{y}_k = \mathbf{A} \, \boldsymbol{\omega}_k \in \mathbb{R}^m$.
- 3. Form orthonormal vectors $\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k \in \mathbb{R}^m$ such that

$$\operatorname{Span}(\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k) = \operatorname{Span}(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k).$$

For instance, Gram-Schmidt can be used — pivoting is rarely required.

If **A** has *exact* rank *k*, then $\text{Span}\{\mathbf{q}_j\}_{j=1}^k = \text{Ran}(\mathbf{A})$ with probability 1.

Range finding problem: Given an $m \times n$ matrix **A** and an integer $k < \min(m, n)$, find an orthonormal $m \times k$ matrix **Q** such that $\mathbf{A} \approx \mathbf{Q}\mathbf{Q}^*\mathbf{A}$.

Solving the primitive problem via randomized sampling — intuition:

- 1. Draw random vectors $\omega_1, \omega_2, \ldots, \omega_k \in \mathbb{R}^n$. (We will discuss the choice of distribution later — think Gaussian for now.)
- 2. Form "sample" vectors $\mathbf{y}_1 = \mathbf{A} \, \boldsymbol{\omega}_1, \, \mathbf{y}_2 = \mathbf{A} \, \boldsymbol{\omega}_2, \, \dots, \, \mathbf{y}_k = \mathbf{A} \, \boldsymbol{\omega}_k \in \mathbb{R}^m$.
- 3. Form orthonormal vectors $\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k \in \mathbb{R}^m$ such that

$$\operatorname{Span}(\mathbf{q}_1, \mathbf{q}_2, \ldots, \mathbf{q}_k) = \operatorname{Span}(\mathbf{y}_1, \mathbf{y}_2, \ldots, \mathbf{y}_k).$$

For instance, Gram-Schmidt can be used — pivoting is rarely required.

If **A** has *exact* rank k, then $\text{Span}\{\mathbf{q}_j\}_{j=1}^k = \text{Ran}(\mathbf{A})$ with probability 1.

What is perhaps surprising is that even in the general case, $\{\mathbf{q}_j\}_{j=1}^k$ often does almost as good of a job as the theoretically optimal vectors (which happen to be the *k* leading left singular vectors).

Range finding problem: Given an $m \times n$ matrix **A** and an integer $k < \min(m, n)$, find an orthonormal $m \times k$ matrix **Q** such that $\mathbf{A} \approx \mathbf{Q}\mathbf{Q}^*\mathbf{A}$.

Solving the primitive problem via randomized sampling — intuition:

- 1. Draw a random matrix $\Omega \in \mathbb{R}^{n \times k}$. (We will discuss the choice of distribution later — think Gaussian for now.)
- 2. Form a "sample" matrix $\mathbf{Y} = \mathbf{A} \, \mathbf{\Omega} \in \mathbb{R}^{m \times k}$.
- 3. Form an orthonormal matrix $\mathbf{Q} \in \mathbb{R}^{m \times k}$ such that $\mathbf{Y} = \mathbf{QR}$. For instance, Gram-Schmidt can be used — pivoting is rarely required.

If **A** has *exact* rank k, then $\mathbf{A} = \mathbf{Q}\mathbf{Q}^*\mathbf{A}$ with probability 1.

Goal: Given an $m \times n$ matrix **A**, compute an approximate rank-k SVD $\mathbf{A} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$.

Algorithm:

- 1. Draw an $n \times k$ Gaussian random matrix $\mathbf{\Omega}$.
- 2. Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.
- 3. Form an $m \times k$ orthonormal matrix **Q** such that **Y** = **Q R**.
- 4. Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
- 5. Compute the SVD of the small matrix **B**: $\mathbf{B} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \mathbf{V}^*$.
- 6. Form the matrix $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.

Goal: Given an $m \times n$ matrix **A**, compute an approximate rank-k SVD **A** \approx **U** Σ **V**^{*}.

Algorithm:

1.	Draw an $n \times k$ Gaussian random matrix $\mathbf{\Omega}$.	Omega = randn(n,1)
2.	Form the $m \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	Y = A * Omega
3.	Form an $m \times k$ orthonormal matrix Q such that Y = Q R .	[Q, R] = qr(Y)
4.	Form the $k \times n$ matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.	B = Q' * A
5.	Compute the SVD of the small matrix B : $\mathbf{B} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \mathbf{V}^*$.	[Uhat, Sigma, V] = svd(B)
6.	Form the matrix $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.	U = Q * Uhat

Single pass algorithms:



Note: With **B** as on the previous slide we have $\mathbf{T} \approx \mathbf{B} \mathbf{Q}$ (sym. case) and $\mathbf{T} \approx \mathbf{B} \mathbf{W}$ (nonsym. case).

Note: An alternative procedure based on "interpolative decompositions" is often more accurate.

References: Woolfe, Liberty, Rokhlin, and Tygert (2008), Clarkson and Woodruff (2009), HMT2009, HMT2011.

Single pass algorithms:

A is symmetric:	A is not symmetric:
Generate a random matrix Ω .	Generate random matrices Ω and Ψ .
Compute a sample matrix Y.	Compute sample matrices $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^* \Psi$.
Find an ON matrix Q	Find ON matrices Q and W
such that $\mathbf{Y} = \mathbf{Q} \mathbf{Q}^* \mathbf{Y}$.	such that $\mathbf{Y} = \mathbf{Q} \mathbf{Q}^* \mathbf{Y}$ and $\mathbf{Z} = \mathbf{W} \mathbf{W}^* \mathbf{Z}$.
Solve for \mathbf{T} the linear system	Solve for T the linear systems
$\mathbf{Q}^* \mathbf{Y} = \mathbf{T} (\mathbf{Q}^* \mathbf{\Omega}).$	$\mathbf{Q}^* \mathbf{Y} = \mathbf{T} \left(\mathbf{W}^* \mathbf{\Omega} ight)$ and $\mathbf{W}^* \mathbf{Z} = \mathbf{T}^* \left(\mathbf{Q}^* \mathbf{\Psi} ight)$.
Factor T so that $\mathbf{T} = \hat{\mathbf{U}} \mathbf{\Lambda} \hat{\mathbf{U}}^*$.	Factor T so that $\mathbf{T} = \hat{\mathbf{U}} \mathbf{\Sigma} \hat{\mathbf{V}}^*$.
Form $\mathbf{U} = \mathbf{Q} \hat{\mathbf{U}}$.	Form $\mathbf{U} = \mathbf{Q} \ \hat{\mathbf{U}}$ and $\mathbf{V} = \mathbf{W} \ \hat{\mathbf{V}}$.
Output: $\mathbf{A} \approx \mathbf{U} \mathbf{\Lambda} \mathbf{U}^*$	Output: $\mathbf{A} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$

Note: With **B** as on the previous slide we have $\mathbf{T} \approx \mathbf{B} \mathbf{Q}$ (sym. case) and $\mathbf{T} \approx \mathbf{B} \mathbf{W}$ (nonsym. case).

Note: An alternative procedure based on "interpolative decompositions" is often more accurate.

References: Woolfe, Liberty, Rokhlin, and Tygert (2008), Clarkson and Woodruff (2009), HMT2009, HMT2011.

Note: If **A** has exact rank k, then the algorithm produces an exact answer with probability 1.

At least three "issues" need attention:

- 1. In practice, **A** does not have exact rank k. There are tail singular modes that pollute the answer.
- 2. In practice, the rank of **A** is not known in advance. Instead, we are typically given a precision ε , and need to find the rank.
- 3. Computations will be carried out in finite precision arithmetic. Since the basis is by nature badly conditioned, care has to be exercised. But:
 - If standard procedures of numerical linear algebra are used (correctly...), things work out fine.
 - However, one must only ask reasonable things of the algorithm: *Good question:* Find U_{approx} , Σ_{approx} , V_{approx} s.t. $||A - U_{approx}\Sigma_{approx}V_{approx}|| \le \varepsilon$. *Bad question:* Find U_{approx} , Σ_{approx} , V_{approx} such that $||U_{exact} - U_{approx}|| \le \varepsilon$, etc.

We will in this talk henceforth ignore question 3 (in the interest of not propagating misconceptions about what numerical analysts find interesting).

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question:

How much does the computation above cost?

Answer:

Cost of steps 2 and 4:Application of A and A* to k vectors.Cost of steps 3,5,6:Dense operations on matrices of sizes $m \times k$ and $k \times n$.

We will consider three proto-typical environments.

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Cost of steps 2 and 4: Application of **A** and \mathbf{A}^* to *k* vectors.

Cost of steps 3,5,6: Dense operations on matrices of sizes $m \times k$ and $k \times n$.

Case 1: A is presented as an array of real numbers in RAM.

Cost is dominated by the 2 m n k flops required for steps 2 and 4.

The O(m n k) flop count is the same as that of standard methods such as Golub-Businger. However, the algorithm above requires no random access to the matrix **A** — the data is retrieved *in two passes.* (One pass in the modified version.)

Remark 1: Improved data access leads to a moderate gain ($\times 2$) in execution time.

Remark 2: Using a special structured random sampling matrix (for instance, a "subsampled random Fourier transform" / "fast Johnson-Lindenstrauss transform"), substantial gain in execution time, and asymptotic complexity of $O(m n \log(k))$ can be achieved.

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Cost of steps 2 and 4: Application of **A** and \mathbf{A}^* to *k* vectors.

Cost of steps 3,5,6: Dense operations on matrices of sizes $m \times k$ and $k \times n$.

Case 1: A is presented as an array of real numbers in RAM.

Use a subsampled random Fourier Transform (SRFT)

 $egin{array}{rcl} m{\Omega} &=& m{D} & m{F} & m{S} \ n imes k & n imes n & n imes n & n imes k \end{array}$

- D is a diagonal matrix whose entries are i.i.d. random variables drawn from a uniform distribution on the unit circle in C.
- **F** is the discrete Fourier transform, $\mathbf{F}_{pq} = \frac{1}{\sqrt{n}} e^{-2\pi i (p-1)(q-1)/n}$.
- S is a matrix whose entries are all zeros except for a single, randomly placed 1 in each column. (In other words, the action of S is to draw k columns at random from DF.)

References: Ailon and Chazelle (2006); Liberty, Rokhlin, Tygert, and Woolfe (2006).

Step (4) involves a second O(mnk) cost, but this can be eliminated via a variation of the randomized sampling process that samples both the column space and the row space:

Given an $m \times n$ matrix **A**, find **U**, **V**, **\Sigma** of rank k such that $\mathbf{A} \approx \mathbf{U} \mathbf{\Sigma} \mathbf{V}^*$.

- 1. Choose an over-sampling parameter p and set $\ell = k + p$. (Say p = k so $\ell = 2k$.)
- 2. Generate SRFT's Ω and Ψ of sizes $n \times \ell$, and $m \times \ell$.
- 3. Form the sample matrices $\mathbf{Y} = \mathbf{A} \, \mathbf{\Omega}$ and $\mathbf{Z} = \mathbf{A}^* \, \Psi$.
- 4. Find ON matrices **Q** and **W** such that $\mathbf{Y} = \mathbf{Q} \mathbf{Q}^* \mathbf{Y}$ and $\mathbf{Z} = \mathbf{W} \mathbf{W}^* \mathbf{Z}$.
- 5. Solve for the $k \times k$ matrix **T** the systems $\mathbf{Q}^* \mathbf{Y} = \mathbf{T} (\mathbf{W}^* \mathbf{\Omega})$ and $\mathbf{W}^* \mathbf{Z} = \mathbf{T}^* (\mathbf{Q}^* \mathbf{\Psi})$.
- 6. Compute the SVD of the small matrix $\mathbf{T} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \hat{\mathbf{V}}^*$ (and truncate if desired).
- 7. Form $\mathbf{U} = \mathbf{Q} \, \hat{\mathbf{U}}$ and $\mathbf{V} = \mathbf{W} \, \hat{\mathbf{V}}$.

Observation 1: Forming **A** Ω and **A**^{*} Ψ in Step 2 has cost $O(m n \log(k))$ since $\ell \sim k$.

Observation 2: All other steps cost at most $O((m+n)k^2)$.

Practical speed of the SRFT-accelerated method

Consider the task of computing a rank-k SVD of a matrix **A** of size $n \times n$.

 $t^{(\text{direct})}$ Time for classical (Golub-Businger) method — $O(k n^2)$

 $t^{(\text{srft})}$ Time for randomized method with an SRFT — $O(\log(k) n^2)$

 $t^{(\text{gauss})}$ Time for randomized method with a Gaussian matrix — $O(k n^2)$

 $t^{(svd)}$ Time for a full SVD — $O(n^3)$

We will show the

acceleration factor =
$$\frac{t^{(\text{direct})}}{t^{(\text{srft})}}$$

for different values of n and k.



Notes:

- Significant speed-ups are achieved for common problem sizes. For instance, $m = n = 2\,000$ and k = 200 leads to a speed-up by roughly a factor of 4.
- Many other choices of random matrices have been found. (Subsampled Hadamard transform; wavelets; random chains of Given's rotations, etc).
- Using a randomly drawn selection of columns as a basis for the range of a matrix **A** is fine, *provided the matrix has first been hit by a random rotation.*
- Current theoretical results seem to overstate the risk of inaccurate results, at least in typical environments.
- The idea was proposed by Liberty, Rokhlin, Tygert, Woolfe (2006).
 - The SRFT (under the name *fast Johnson-Lindenstrauss transform*) was suggested by Nir Ailon and Bernard Chazelle (2006) in a related context.
 - Dissertation by Edo Liberty.
 - Related recent work by Sarlós (on randomized regression).
 - Overdetermined linear least-squares regression: Rokhlin and Tygert (2008).

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Cost of steps 2 and 4: Application of **A** and \mathbf{A}^* to *k* vectors.

Cost of steps 3,5,6: Dense operations on matrices of sizes $m \times k$ and $k \times n$.

Case 2: A is presented as an array of real numbers in slow memory — on "disk".

In this case, standard methods such as Golub-Businger become prohibitively slow due to the random access requirements on the data.

However, the method described above works just fine. (Recall: Two passes only!)

Limitation 1: Matrices of size $m \times k$ and $k \times n$ must fit in RAM.

Limitation 2: For matrices whose singular values decay slowly (as is typical in the data-analysis environment), the method above is typically not accurate enough. *Higher accuracy can be bought at modest cost — we will return to this point.*

Code: Webpage of Yoel Shkolnisky.



Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Cost of steps 2 and 4: Application of **A** and \mathbf{A}^* to *k* vectors.

Cost of steps 3,5,6: Dense operations on matrices of sizes $m \times k$ and $k \times n$.

Case 3: A and A* admit fast matrix-vector multiplication.

In this case, the "standard" method is some variation of Krylov methods such as Lanczos (or Arnoldi for non-symmetric matrices) whose cost T_{Krylov} satisfy:

$$T_{\rm Krylov} \sim k T_{\rm mat}$$
-vec-mult + $\underbrace{O(k^2 (m+n))}_{\rm With \ full \ "reorthogonalization"}$.

The asymptotic cost of the randomized scheme is the same; its advantage is again in how the data is accessed — the k matrix-vector multiplies can be executed in parallel.

The method above is in important environments less accurate than Arnoldi, but this can be fixed while compromising only slightly on the pass-efficiency.

THEORY

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = ||\mathbf{A} - \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*||$? (Recall $e_k = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||$.)

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = ||\mathbf{A} - \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*||$? (Recall $e_k = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||$.)

Eckart-Young theorem: e_k is bounded from below by the (k+1)'th singular value σ_{k+1} of **A**.

Question: Is e_k close to σ_{k+1} ?

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = ||\mathbf{A} - \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*||$? (Recall $e_k = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||$.)

Eckart-Young theorem: e_k is bounded from below by the (k+1)'th singular value σ_{k+1} of **A**.

Question: Is e_k close to σ_{k+1} ?

Answer: Lamentably, no. The expectation of $\frac{e_k}{\sigma_{k+1}}$ is large, and has very large variance.

Output: Rank-*k* factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

(1) Draw an $n \times k$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times k$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

Question: What is the error $e_k = ||\mathbf{A} - \mathbf{U}\mathbf{\Sigma}\mathbf{V}^*||$? (Recall $e_k = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||$.)

Eckart-Young theorem: e_k is bounded from below by the (k+1)'th singular value σ_{k+1} of **A**.

Question: Is e_k close to σ_{k+1} ?

Answer: Lamentably, no. The expectation of $\frac{e_k}{\sigma_{k+1}}$ is large, and has very large variance.

Remedy: Over-sample slightly. Compute k+p samples from the range of **A**. It turns out that p = 5 or 10 is often sufficient. p = k is almost always more than enough.

Input: An $m \times n$ matrix A , a target rank k, and an over-sampling parameter p (say $p = 5$).		
<i>Output:</i> Rank- $(k + p)$ factors U , S , and V in an approximate SVD A \approx USV [*] .		
(1) Draw an $n \times (k+p)$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.	
(2) Form the $n \times (k + p)$ sample matrix $\mathbf{Y} = \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \mathbf{\Sigma} \mathbf{V}^*$.	
(3) Compute an ON matrix Q s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.	
Bound on the expectation of the error for Gaussian test matrices

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let $\boldsymbol{\Omega}$ denote an $n \times (k + p)$ Gaussian matrix.

Let \mathbf{Q} denote the $m \times (k + p)$ matrix $\mathbf{Q} = \operatorname{orth}(\mathbf{A}\mathbf{\Omega})$.

If $p \ge 2$, then

$$\mathbb{E}||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||_{\text{Frob}} \le \left(1 + \frac{k}{p-1}\right)^{1/2} \left(\sum_{\substack{j=k+1 \\ j=k+1}}^{\min(m,n)} \sigma_j^2\right)^{1/2},$$

and

$$\mathbb{E}||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{\substack{j=k+1}}^{\min(m,n)}\sigma_j^2\right)^{1/2}$$

Ref: Halko, Martinsson, Tropp, 2009 & 2011

Large deviation bound for the error for Gaussian test matrices

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let $\boldsymbol{\Omega}$ denote an $n \times (k + p)$ Gaussian matrix.

Let **Q** denote the $m \times (k + p)$ matrix **Q** = orth(**A** Ω).

If $p \ge 4$, and u and t are such that $u \ge 1$ and $t \ge 1$, then

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + t\sqrt{\frac{3k}{p+1}} + ut\frac{e\sqrt{k+p}}{p+1}\right)\sigma_{k+1} + \frac{te\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}$$

1 10

10

except with probability at most $2t^{-p} + e^{-u^2/2}$.

Ref: Halko, Martinsson, Tropp, 2009 & 2011; Martinsson, Rokhlin, Tygert (2006)

u and t parameterize "bad" events — large u, t is bad, but unlikely.

Certain choices of t and u lead to simpler results. For instance,

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + 16\sqrt{1 + \frac{k}{p+1}}\right)\sigma_{k+1} + 8\frac{\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2},$$

except with probability at most $3 e^{-p}$.

Large deviation bound for the error for Gaussian test matrices

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter.

Let $\boldsymbol{\Omega}$ denote an $n \times (k + p)$ Gaussian matrix.

Let \mathbf{Q} denote the $m \times (k + p)$ matrix $\mathbf{Q} = \operatorname{orth}(\mathbf{A}\mathbf{\Omega})$.

If $p \ge 4$, and u and t are such that $u \ge 1$ and $t \ge 1$, then

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + t\sqrt{\frac{3k}{p+1}} + ut\frac{e\sqrt{k+p}}{p+1}\right)\sigma_{k+1} + \frac{te\sqrt{k+p}}{p+1}\left(\sum_{j>k}\sigma_j^2\right)^{1/2}$$

10

except with probability at most $2t^{-p} + e^{-u^2/2}$.

Ref: Halko, Martinsson, Tropp, 2009 & 2011; Martinsson, Rokhlin, Tygert (2006)

u and t parameterize "bad" events — large u, t is bad, but unlikely.

Certain choices of t and u lead to simpler results. For instance,

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + 6\sqrt{(k+p) \cdot p \log p}\right)\sigma_{k+1} + 3\sqrt{k+p}\left(\sum_{j>k}\sigma_j^2\right)^{1/2},$$

except with probability at most $3 p^{-p}$.

Proofs — Overview:

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.

Let Ω denote an $n \times \ell$ "test matrix", and let \mathbf{Q} denote the $m \times \ell$ matrix $\mathbf{Q} = \texttt{orth}(\mathbf{A}\Omega)$.

We seek to bound the error $e_k = e_k(\mathbf{A}, \mathbf{\Omega}) = ||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||$, which is a random variable.

1. Make no assumption on Ω . Construct a deterministic bound of the form

 $||\textbf{A} - \textbf{Q}\textbf{Q}^*\textbf{A}|| \leq \cdots \textbf{A} \cdots \boldsymbol{\Omega} \cdots$

2. Assume that Ω is drawn from a normal Gaussian distribution. Take expectations of the deterministic bound to attain a bound of the form

$$\mathbb{E}\big[||\textbf{A} - \textbf{Q}\textbf{Q}^*\textbf{A}||\big] \leq \cdots \textbf{A} \cdots$$

3. Assume that Ω is drawn from a normal Gaussian distribution. Take expectations of the deterministic bound conditioned on "bad behavior" in Ω to get

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \leq \cdots \mathbf{A} \cdots$$

that hold with probability at least \cdots .

Part 1 (out of 3) — deterministic bound:

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.

Let Ω denote an $n \times \ell$ "test matrix", and let \mathbf{Q} denote the $m \times \ell$ matrix $\mathbf{Q} = orth(\mathbf{A}\Omega)$.

Partition the SVD of **A** as follows:

$$\mathbf{A} = \mathbf{U} \quad \begin{bmatrix} \mathbf{x}_1 & \\ & \mathbf{x}_2 \end{bmatrix} \begin{bmatrix} \mathbf{v}_1^* \\ & \mathbf{v}_2^* \end{bmatrix} \quad \begin{array}{c} k \\ & n-k \end{array}$$

Define Ω_1 and Ω_2 via

$$\boldsymbol{\Omega}_1 = \boldsymbol{\mathsf{V}}_1^* \, \boldsymbol{\Omega} \quad \text{and} \quad \boldsymbol{\Omega}_2 = \boldsymbol{\mathsf{V}}_2^* \, \boldsymbol{\Omega}.$$

Theorem: [HMT2009,HMT2011] Assuming that Ω_1 is not singular, it holds that

$$|||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|||^2 \leq \underbrace{|||\mathbf{\Sigma}_2|||^2}_{\text{theoretically minimal error}} + |||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}|||^2$$

Here, $||| \cdot |||$ represents either ℓ^2 -operator norm, or the Frobenius norm.

Note: A similar (but weaker) result appears in Boutsidis, Mahoney, Drineas (2009).

Recall:
$$\mathbf{A} = \mathbf{U} \begin{bmatrix} \mathbf{\Sigma}_1 & \mathbf{0} \\ \mathbf{0} & \mathbf{\Sigma}_2 \end{bmatrix} \begin{bmatrix} \mathbf{V}_1^* \\ \mathbf{V}_2^* \end{bmatrix}, \begin{bmatrix} \mathbf{\Omega}_1 \\ \mathbf{\Omega}_2 \end{bmatrix} = \begin{bmatrix} \mathbf{V}_1^* \mathbf{\Omega} \\ \mathbf{V}_2^* \mathbf{\Omega} \end{bmatrix}, \mathbf{Y} = \mathbf{A}\mathbf{\Omega}, \mathbf{P} \text{ proj}^n \text{ onto } \text{Ran}(\mathbf{Y}).$$

Thm: Suppose $\Sigma_1 \Omega_1$ has full rank. Then $||\mathbf{A} - \mathbf{PA}||^2 \le ||\Sigma_2||^2 + ||\Sigma_2 \Omega_2 \Omega_1^{\dagger}||^2$.

Proof: The problem is rotationally invariant \Rightarrow We can assume U = I and so A = ΣV^* . Simple calculation: $||(\mathbf{I} - \mathbf{P})\mathbf{A}||^2 = ||\mathbf{A}^*(\mathbf{I} - \mathbf{P})^2\mathbf{A}|| = ||\mathbf{\Sigma}(\mathbf{I} - \mathbf{P})\mathbf{\Sigma}||.$ $\operatorname{Ran}(\mathbf{Y}) = \operatorname{Ran}\left(\left| \begin{array}{c} \mathbf{\Sigma}_{1} \mathbf{\Omega}_{1} \\ \mathbf{\Sigma}_{2} \mathbf{\Omega}_{2} \end{array} \right| \right) = \operatorname{Ran}\left(\left| \begin{array}{c} \mathbf{I} \\ \mathbf{\Sigma}_{2} \mathbf{\Omega}_{2} \mathbf{\Omega}_{1}^{\dagger} \mathbf{\Sigma}_{1} \end{array} \right| \mathbf{\Sigma}_{1} \mathbf{\Omega}_{1} \right) = \operatorname{Ran}\left(\left| \begin{array}{c} \mathbf{I} \\ \mathbf{\Sigma}_{2} \mathbf{\Omega}_{2} \mathbf{\Omega}_{1}^{\dagger} \mathbf{\Sigma}_{1} \end{array} \right| \right)$ Set $\mathbf{F} = \mathbf{\Sigma}_2 \mathbf{\Omega}_2 \mathbf{\Omega}_1^{\dagger} \mathbf{\Sigma}_1^{-1}$. Then $\mathbf{P} = \begin{bmatrix} \mathbf{I} \\ \mathbf{F} \end{bmatrix} (\mathbf{I} + \mathbf{F}^* \mathbf{F})^{-1} [\mathbf{I} \mathbf{F}^*]$. (Compare to $\mathbf{P}_{\text{ideal}} = \begin{bmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$.) Use properties of psd matrices: $\mathbf{I} - \mathbf{P} \preccurlyeq \cdots \preccurlyeq \begin{bmatrix} \mathbf{F}^* \mathbf{F} & -(\mathbf{I} + \mathbf{F}^* \mathbf{F})^{-1} \mathbf{F}^* \\ -\mathbf{F}(\mathbf{I} + \mathbf{F}^* \mathbf{F})^{-1} & \mathbf{I} \end{bmatrix}$ $\begin{array}{c} \text{Conjugate by } \boldsymbol{\Sigma} \text{ to get } \boldsymbol{\Sigma}(\boldsymbol{\mathsf{I}}-\boldsymbol{\mathsf{P}})\boldsymbol{\Sigma} \preccurlyeq \left[\begin{array}{c} \boldsymbol{\Sigma}_1\boldsymbol{\mathsf{F}}^*\boldsymbol{\mathsf{F}}\boldsymbol{\Sigma}_1 & -\boldsymbol{\Sigma}_1(\boldsymbol{\mathsf{I}}+\boldsymbol{\mathsf{F}}^*\boldsymbol{\mathsf{F}})^{-1}\boldsymbol{\mathsf{F}}^*\boldsymbol{\Sigma}_2 \\ -\boldsymbol{\Sigma}_2\boldsymbol{\mathsf{F}}(\boldsymbol{\mathsf{I}}+\boldsymbol{\mathsf{F}}^*\boldsymbol{\mathsf{F}})^{-1}\boldsymbol{\Sigma}_1 & \boldsymbol{\Sigma}_2^2 \end{array} \right] \end{array}$

 $\text{Diagonal dominance: } ||\boldsymbol{\Sigma}(\boldsymbol{\mathsf{I}}-\boldsymbol{\mathsf{P}})\boldsymbol{\Sigma}|| \leq ||\boldsymbol{\Sigma}_1\boldsymbol{\mathsf{F}}^*\boldsymbol{\mathsf{F}}\boldsymbol{\Sigma}_1|| + ||\boldsymbol{\Sigma}_2^2|| = ||\boldsymbol{\Sigma}_2\boldsymbol{\Omega}_2\boldsymbol{\Omega}_1^\dagger||^2 + ||\boldsymbol{\Sigma}_2||^2.$

Part 2 (out of 3) — bound on expectation of error when Ω is Gaussian:

Let **A** denote an $m \times n$ matrix with singular values $\{\sigma_j\}_{j=1}^{\min(m,n)}$.

Let k denote a target rank and let p denote an over-sampling parameter. Set $\ell = k + p$.

Let Ω denote an $n \times \ell$ "test matrix", and let \mathbf{Q} denote the $m \times \ell$ matrix $\mathbf{Q} = orth(\mathbf{A}\Omega)$.

Recall: $|||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|||^2 \le |||\mathbf{\Sigma}_2|||^2 + |||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}|||^2$, where $\mathbf{\Omega}_1 = \mathbf{V}_1^*\mathbf{\Omega}$ and $\mathbf{\Omega}_2 = \mathbf{V}_2^*\mathbf{\Omega}$.

Assumption: Ω is drawn from a normal Gaussian distribution.

Since the Gaussian distribution is rotationally invariant, the matrices Ω_1 and Ω_2 also have a Gaussian distribution. (As a consequence, the matrices U and V do not enter the analysis and one could simply assume that A is diagonal, $A = \text{diag}(\sigma_1, \sigma_2, ...)$.)

What is the distribution of Ω_1^{\dagger} when Ω_1 is a $k \times (k + p)$ Gaussian matrix?

If p = 0, then $||\mathbf{\Omega}_1^{\dagger}||$ is typically large, and is very unstable.



Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. p = 0



Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. p = 2



Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. p = 5



Scatter plot showing distribution of $1/\sigma_{\min}$ for $k \times (k+p)$ Gaussian matrices. p = 10



Scatter plot showing distribution of $k \times (k + p)$ Gaussian matrices.

Simplistic proof that a rectangular Gaussian matrix is well-conditioned:

Let **G** denote a $k \times \ell$ Gaussian matrix where $k < \ell$. Let "g" denote a generic $\mathcal{N}(0,1)$ variable and " r_j^2 " denote a generic χ_j^2 variable. Then

$$\mathbf{G} \sim \begin{bmatrix} g & g & g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & g & \cdots \\ g & g & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \end{bmatrix} \sim \begin{bmatrix} r_{\ell} & 0 & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & g & g & g & g & g & g & \cdots \\ 0 & g & g & g & g & g & g & \cdots \\ 0 & g & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \end{bmatrix} \sim \begin{bmatrix} r_{\ell} & 0 & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & g & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \end{bmatrix} \sim \begin{bmatrix} r_{\ell} & 0 & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & g & g & g & g & g & \cdots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \ddots \end{bmatrix} \sim \begin{bmatrix} r_{\ell} & 0 & 0 & 0 & 0 & 0 & \cdots \\ r_{k-1} & r_{\ell-1} & 0 & 0 & 0 & \cdots \\ 0 & r_{k-2} & g & g & g & \cdots \\ 0 & 0 & r_{k-2} & r_{\ell-2} & 0 & 0 & \cdots \\ 0 & 0 & r_{k-3} & r_{\ell-3} & 0 & \cdots \\ \vdots & \ddots & \ddots \end{bmatrix}$$

Gershgorin's circle theorem will now show that **G** is well-conditioned if, e.g., $\ell = 2k$. More sophisticated methods are required to get to $\ell = k + 2$. **Some results on Gaussian matrices.** Adapted from HMT2009/2011; Gordon (1985,1988) for Proposition 1; Chen & Dongarra (2005) for Propositions 2 and 4; Bogdanov (1998) for Proposition 3.

Proposition 1: Let G be a Gaussian matrix. Then

(1)
$$\left(\mathbb{E}\left[||\mathbf{SGT}||_{\mathrm{F}}^{2}
ight]
ight)^{1/2} \leq ||\mathbf{S}||_{\mathrm{F}} \, ||\mathbf{T}||_{\mathrm{F}}$$

(2)
$$\mathbb{E}[||\mathbf{SGT}||] \leq ||\mathbf{S}|| \, ||\mathbf{T}||_{\mathrm{F}} + ||\mathbf{S}||_{\mathrm{F}} \, ||\mathbf{T}||$$

Proposition 2: Let **G** be a Gaussian matrix of size $k \times k + p$ where $p \ge 2$. Then

(3)

$$\left(\mathbb{E}\left[||\mathbf{G}^{\dagger}||_{\mathrm{F}}^{2}\right]\right)^{1/2} \leq \sqrt{\frac{k}{p-1}}$$
(4)

$$\mathbb{E}\left[||\mathbf{G}^{\dagger}||\right] \leq \frac{e\sqrt{k+p}}{p}.$$

Proposition 3: Suppose *h* is Lipschitz $|h(\mathbf{X}) - h(\mathbf{Y})| \le L||\mathbf{X} - \mathbf{Y}||_{\mathrm{F}}$ and **G** is Gaussian. Then

(5)
$$\mathbb{P}[h(\mathbf{G}) > \mathbb{E}[h(\mathbf{G})] + L u] \le e^{-u^2/2}.$$

Proposition 4: Suppose **G** is Gaussian of size $k \times k + p$ with $p \ge 4$. Then for $t \ge 1$:

(6)

$$\mathbb{P}\left[||\mathbf{G}^{\dagger}||_{\mathrm{F}} \geq \sqrt{\frac{3k}{p+1}}t\right] \leq t^{-p}$$
(7)

$$\mathbb{P}\left[||\mathbf{G}^{\dagger}|| \geq \frac{e\sqrt{k+p}}{p+1}t\right] \leq t^{-(p+1)}$$

Recall: $||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||^2 \le ||\mathbf{\Sigma}_2||^2 + ||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}||^2$, where $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$ are Gaussian and $\mathbf{\Omega}_1$ is $k \times k + p$.

Theorem:
$$\mathbb{E}\left[||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||\right] \le \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j=k+1}^{\min(m,n)}\sigma_j^2\right)^{1/2}$$
.

Proof: First observe that

 $\mathbb{E}||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| = \mathbb{E}(||\mathbf{\Sigma}_2||^2 + ||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}||^2)^{1/2} \le ||\mathbf{\Sigma}_2|| + \mathbb{E}||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}||.$

Condition on Ω_1 and use Proposition 1:

$$\begin{split} \mathbb{E}||\boldsymbol{\Sigma}_{2}\boldsymbol{\Omega}_{2}\boldsymbol{\Omega}_{1}^{\dagger}|| &\leq \mathbb{E}\big[||\boldsymbol{\Sigma}_{2}|| \, ||\boldsymbol{\Omega}_{1}^{\dagger}||_{\mathrm{F}} + ||\boldsymbol{\Sigma}_{2}||_{\mathrm{F}} \, ||\boldsymbol{\Omega}_{1}^{\dagger}||\big] \\ &\leq \{\mathsf{H\"{o}lder}\} \leq ||\boldsymbol{\Sigma}_{2}|| \, \big(\mathbb{E}||\boldsymbol{\Omega}_{1}^{\dagger}||_{\mathrm{F}}^{2}\big)^{1/2} + ||\boldsymbol{\Sigma}_{2}||_{\mathrm{F}} \, \mathbb{E}||\boldsymbol{\Omega}_{1}^{\dagger}||. \end{split}$$

Proposition 2 now provides bounds for $\mathbb{E}||\Omega_1^{\dagger}||_F^2$ and $\mathbb{E}||\Omega_1^{\dagger}||$ and we get

$$\mathbb{E}||\boldsymbol{\Sigma}_{2}\boldsymbol{\Omega}_{2}\boldsymbol{\Omega}_{1}^{\dagger}|| \leq \sqrt{\frac{k}{p-1}}||\boldsymbol{\Sigma}_{2}|| + \frac{e\sqrt{k+p}}{p}||\boldsymbol{\Sigma}_{2}||_{\mathrm{F}} = \sqrt{\frac{k}{p-1}}\sigma_{k+1} + \frac{e\sqrt{k+p}}{p}\left(\sum_{j>k}\sigma_{j}^{2}\right)^{1/2}$$

Some results on Gaussian matrices. Adapted from HMT2009/2011; Gordon (1985,1988) for Proposition 1; Chen & Dongarra (2005) for Propositions 2 and 4; Bogdanov (1998) for Proposition 3.

Proposition 1: Let G be a Gaussian matrix. Then

(8)
$$\left(\mathbb{E}\left[||\mathbf{SGT}||_{\mathrm{F}}^{2}\right]\right)^{1/2} \leq ||\mathbf{S}||_{\mathrm{F}} \, ||\mathbf{T}||_{\mathrm{F}}$$

(9)
$$\mathbb{E}[||\mathbf{SGT}||] \leq ||\mathbf{S}|| \, ||\mathbf{T}||_{\mathrm{F}} + ||\mathbf{S}||_{\mathrm{F}} \, ||\mathbf{T}||_{\mathrm{F}}$$

Proposition 2: Let **G** be a Gaussian matrix of size $k \times k + p$ where $p \ge 2$. Then

(10)

$$\left(\mathbb{E}\left[||\mathbf{G}^{\dagger}||_{\mathrm{F}}^{2}\right]\right)^{1/2} \leq \sqrt{\frac{k}{p-1}}$$

$$\mathbb{E}\left[||\mathbf{G}^{\dagger}||\right] \leq \frac{e\sqrt{k+p}}{p}.$$

Proposition 3: Suppose *h* is Lipschitz $|h(\mathbf{X}) - h(\mathbf{Y})| \le L||\mathbf{X} - \mathbf{Y}||_{\mathrm{F}}$ and **G** is Gaussian. Then

(12)
$$\mathbb{P}[h(\mathbf{G}) > \mathbb{E}[h(\mathbf{G})] + L u] \le e^{-u^2/2}.$$

Proposition 4: Suppose **G** is Gaussian of size $k \times k + p$ with $p \ge 4$. Then for $t \ge 1$:

(13)
$$\mathbb{P}\left[||\mathbf{G}^{\dagger}||_{\mathrm{F}} \ge \sqrt{\frac{3k}{p+1}}t\right] \le t^{-p}$$

(14)
$$\mathbb{P}\left[||\mathbf{G}^{\dagger}|| \ge \frac{e\sqrt{k+p}}{p+1}t\right] \le t^{-(p+1)}$$

Recall: $||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}||^2 \le ||\mathbf{\Sigma}_2||^2 + ||\mathbf{\Sigma}_2\mathbf{\Omega}_2\mathbf{\Omega}_1^{\dagger}||^2$, where $\mathbf{\Omega}_1$ and $\mathbf{\Omega}_2$ are Gaussian and $\mathbf{\Omega}_1$ is $k \times k + p$.

Theorem: With probability at least $1 - 2t^{-p} - e^{-u^2/2}$ it holds that

$$||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + t\sqrt{\frac{3k}{p+1}} + ut\frac{e\sqrt{k+p}}{p+1}\right) \sigma_{k+1} + \frac{te\sqrt{k+p}}{p+1} \left(\sum_{j>k}\sigma_j^2\right)^{1/2}$$

Proof: Set $E_t = \left\{ ||\mathbf{\Omega}_1|| \le \frac{e\sqrt{k+p}}{p+1}t \text{ and } ||\mathbf{\Omega}_1^{\dagger}||_{\mathrm{F}} \le \sqrt{\frac{3k}{p+1}}t \right\}$. By Proposition 4: $\mathbb{P}(E_t^{\mathrm{c}}) \le 2t^{-p}$.

Set $h(\mathbf{X}) = ||\mathbf{\Sigma}_2 \mathbf{X} \mathbf{\Omega}_1^{\dagger}||$. A direct calculation shows

$$|h(\mathbf{X}) - h(\mathbf{Y})| \le ||\mathbf{\Sigma}_2|| \, ||\mathbf{\Omega}_1^{\dagger}|| \, ||\mathbf{X} - \mathbf{y}||_{\mathrm{F}}.$$

Hold Ω_1 fixed and take the expectation on Ω_2 . Then Proposition 1 applies and so

$$\mathbb{E} \big[h \big(\boldsymbol{\Omega}_2 \big) \mid \boldsymbol{\Omega}_1 \big] \leq || \boldsymbol{\Sigma}_2 || \, || \boldsymbol{\Omega}_1^{\dagger} ||_{\mathrm{F}} + || \boldsymbol{\Sigma}_2 ||_{\mathrm{F}} \, || \boldsymbol{\Omega}_1^{\dagger} ||.$$

Now use Proposition 3 (concentration of measure)

$$\mathbb{P}\left[\underbrace{||\boldsymbol{\Sigma}_{2}\boldsymbol{\Omega}_{2}\boldsymbol{\Omega}_{1}^{\dagger}||}_{=h(\boldsymbol{\Omega}_{2})} > \underbrace{||\boldsymbol{\Sigma}_{2}|| \, ||\boldsymbol{\Omega}_{1}^{\dagger}||_{\mathrm{F}} + ||\boldsymbol{\Sigma}_{2}||_{\mathrm{F}} \, ||\boldsymbol{\Omega}_{1}^{\dagger}||}_{=\mathbb{E}[h(\boldsymbol{\Omega}_{2})]} + \underbrace{||\boldsymbol{\Sigma}_{2}|| \, ||\boldsymbol{\Omega}_{1}^{\dagger}||}_{=L} \, u \mid E_{t}\right] < e^{-u^{2}/2}.$$

When E_t holds true, we have bounds on the "badness" of Ω_1^{\dagger} :

$$\mathbb{P}\Big[||\mathbf{\Sigma}_{2}\mathbf{\Omega}_{2}\mathbf{\Omega}_{1}^{\dagger}|| > ||\mathbf{\Sigma}_{2}|| \sqrt{\frac{3k}{p+1}}t + ||\mathbf{\Sigma}_{2}||_{\mathrm{F}} \frac{e\sqrt{k+p}}{p+1}t + ||\mathbf{\Sigma}_{2}|| \frac{e\sqrt{k+p}}{p+1}ut \mid E_{t}\Big] < e^{-u^{2}/2}.$$

The theorem is obtained by using $\mathbb{P}(E_t^c) \leq 2t^{-p}$ to remove the conditioning of E_t .

We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:



We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:



We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:



We consider a $1\,000 \times 1\,000$ matrix **A** whose singular values are shown below:



Example 2:

The matrix **A** being analyzed is a 9025×9025 matrix arising in a diffusion geometry approach to image processing.

To be precise, **A** is a graph Laplacian on the manifold of 3×3 patches.



Joint work with François Meyer of the University of Colorado at Boulder.



The pink lines illustrates the performance of the basic random sampling scheme. The errors are huge, and the estimated eigenvalues are much too small.

Example 3: "Eigenfaces"

We next process a data base containing m = 7254 pictures of faces

Each image consists of $n = 384 \times 256 = 98304$ gray scale pixels.

We center and scale the pixels in each image, and let the resulting values form a column of a $98\,304 \times 7\,254$ data matrix **A**.

The left singular vectors of **A** are the so called *eigenfaces* of the data base.



The pink lines illustrates the performance of the basic random sampling scheme. Again, the errors are huge, and the estimated eigenvalues are much too small.

Power method for improving accuracy:

The error depends on how quickly the singular values decay. Recall that

$$\mathbb{E}||\mathbf{A} - \mathbf{Q}\mathbf{Q}^*\mathbf{A}|| \le \left(1 + \sqrt{\frac{k}{p-1}}\right)\sigma_{k+1} + \frac{e\sqrt{k+p}}{p} \left(\sum_{\substack{j=k+1}}^{\min(m,n)} \sigma_j^2\right)^{1/2}$$

The faster the singular values decay — the stronger the relative weight of the dominant modes in the samples.

Idea: The matrix $(\mathbf{A} \mathbf{A}^*)^q \mathbf{A}$ has the same left singular vectors as \mathbf{A} , and singular values

$$\sigma_j((\mathbf{A} \mathbf{A}^*)^q \mathbf{A}) = (\sigma_j(\mathbf{A}))^{2q+1}.$$

Much faster decay — so let us use the sample matrix

$$\mathbf{Y} = (\mathbf{A} \, \mathbf{A}^*)^q \, \mathbf{A} \, \mathbf{\Omega}$$

instead of

$$\mathbf{Y} = \mathbf{A} \, \mathbf{\Omega}.$$

References: Paper by Rokhlin, Szlam, Tygert (2008). Suggestions by Ming Gu. Also similar to "block power method," "block Lanczos," "subspace iteration."

Input: An $m \times n$ matrix **A**, a target rank ℓ , and a small integer q.

Output: Rank- ℓ factors U, Σ , and V in an approximate SVD $\mathbf{A} \approx \mathbf{U} \Sigma \mathbf{V}^*$.

(1) Draw an $n \times \ell$ random matrix $\mathbf{\Omega}$.	(4) Form the small matrix $\mathbf{B} = \mathbf{Q}^* \mathbf{A}$.
(2) Form the $n \times \ell$ sample matrix $\mathbf{Y} = (\mathbf{A} \mathbf{A}^*)^q \mathbf{A} \mathbf{\Omega}$.	(5) Factor the small matrix $\mathbf{B} = \hat{\mathbf{U}} \boldsymbol{\Sigma} \mathbf{V}^*$.
(3) Compute an ON matrix \mathbf{Q} s.t. $\mathbf{Y} = \mathbf{Q}\mathbf{Q}^*\mathbf{Y}$.	(6) Form $\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}$.

• Detailed (and, we believe, close to sharp) error bounds have been proven. For instance, the expectation of the error satisfies:

(15)
$$\mathbb{E}\left[||\mathbf{A} - \mathbf{A}_{k+p}^{\text{computed}}||\right] \le \left(1 + 4\sqrt{\frac{2\min(m,n)}{k-1}}\right)^{1/(2q+1)} \sigma_{k+1}(\mathbf{A}).$$

Reference: Halko, Martinsson, Tropp (2011).

- The improved accuracy from the modified scheme comes at a cost; 2q+1 passes over the matrix are required instead of 1. However, q can often be chosen quite small in practice, q = 2 or q = 3, say.
- The bound (15) assumes exact arithmetic.
 To handle round-off errors, variations of subspace iterations can be used.
 These are entirely numerically stable and achieve the same error bound.

A numerically stable version of the "power method":

Input: An $m \times n$ matrix **A**, a target rank ℓ , and a small integer q. *Output:* Rank- ℓ factors **U**, Σ , and **V** in an approximate SVD **A** \approx **U** Σ **V**^{*}.

```
Draw an n \times \ell Gaussian random matrix \Omega.
Set \mathbf{Q} = \operatorname{orth}(\mathbf{A}\Omega)
for i = 1, 2, \dots, q
\mathbf{W} = \operatorname{orth}(\mathbf{A}^*\mathbf{Q})
\mathbf{Q} = \operatorname{orth}(\mathbf{A}\mathbf{W})
end for
\mathbf{B} = \mathbf{Q}^*\mathbf{A}
[\hat{\mathbf{U}}, \Sigma, \mathbf{V}] = \operatorname{svd}(\mathbf{B})
\mathbf{U} = \mathbf{Q}\hat{\mathbf{U}}.
```

Note: Algebraically, the method with orthogonalizations is identical to the "original" method where $\mathbf{Q} = \operatorname{orth}((\mathbf{A}\mathbf{A}^*)^q \mathbf{A}\Omega)$.

Note: This is a classic subspace iteration.

The novelty is the error analysis, and the finding that using a very small q is often fine. (In fact, our analysis allows q to be zero...)

Example 2 (revisited):

The matrix **A** being analyzed is a 9025×9025 matrix arising in a diffusion geometry approach to image processing.

To be precise, **A** is a graph Laplacian on the manifold of 3×3 patches.



Joint work with François Meyer of the University of Colorado at Boulder.



The pink lines illustrates the performance of the basic random sampling scheme. The errors are huge, and the estimated eigenvalues are much too small.

Example 3 (revisited): "Eigenfaces"

We next process a data base containing m = 7254 pictures of faces

Each image consists of $n = 384 \times 256 = 98304$ gray scale pixels.

We center and scale the pixels in each image, and let the resulting values form a column of a $98\,304 \times 7\,254$ data matrix **A**.

The left singular vectors of **A** are the so called *eigenfaces* of the data base.



The pink lines illustrates the performance of the basic random sampling scheme. Again, the errors are huge, and the estimated eigenvalues are much too small.

Example of application: Nearest neighbor search in \mathbb{R}^D

Peter Jones, Andrei Osipov, Vladimir Rokhlin

Suppose you are given N points $\{x_j\}_{j=1}^N$ in \mathbb{R}^D . The coordinate matrix is

$$\mathbf{X} = [oldsymbol{x}_1 \; oldsymbol{x}_2 \; \cdots oldsymbol{x}_N] \in \mathbb{R}^{D imes N}$$

How do you find the k nearest neighbors for every point?

If D = 2 there are several options; you can, e.g, sort the points into a quad-tree. Standard techniques get very expensive as D grows.

Simple idea: Use a random map to project onto low-dimensional space. This "sort of" preserves distances. Execute a fast search there.

Improved idea: The output from a single random projection is unreliable. But, you can repeat the experiment several times, use these to generate a list of *candidates* for the nearest neighbors, and then compute exact distances to find the k closest among the candidates.

Final remarks:

- For large scale SVD/PCA of dense matrices, these algorithms are highly recommended; they compare favorably to existing methods in almost every regard. Free software can be downloaded → google *Mark Tygert*.
- The approximation error is a random variable, but its distribution is very narrowly concentrated. Rigorous error bounds that are satisfied with probability 1η where η is a user set "failure probability" (*e.g.* $\eta = 10^{-10}$ or 10^{-20}).
- This talk mentioned *error estimators* only briefly, but they are important. Can operate independently of the algorithm for improved robustness. Typically cheap and easy to implement. Used to determine the actual rank.
- The theory can be hard, but *experimentation is easy!* Concentration of measure makes the algorithms behave as if deterministic.
- Randomized methods for computing "FMM"-style (HSS, *H*-matrix, ...) representations of matrices have been developed [Martinsson 2008 & 2011], [Lin, Lu, Ying 2011].
- Survey paper: N. Halko, P.G. Martinsson, J. Tropp, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions.* SIAM Review, 53(2), 2011, pp. 217–288.
 Slides: Under "Talks" tab on my webpage (google "Gunnar Martinsson").

We note that the power method can be viewed as a hybrid between the "basic" randomized method, and a Krylov subspace method:

Krylov method: Restrict A to the linear space

$$\mathcal{V}_q(\boldsymbol{\omega}) = \operatorname{Span}(\mathbf{A}\boldsymbol{\omega}, \ \mathbf{A}^2\boldsymbol{\omega}, \dots, \ \mathbf{A}^q\boldsymbol{\omega}).$$

"Basic" randomized method: Restrict A to the linear space

$$\mathsf{Span}(\mathsf{A}\boldsymbol{\omega}_1, \, \mathsf{A}\boldsymbol{\omega}_2, \, \dots, \, \mathsf{A}\boldsymbol{\omega}_\ell) = \mathcal{V}_1(\boldsymbol{\omega}_1) \times \mathcal{V}_1(\boldsymbol{\omega}_2) \times \dots \times \mathcal{V}_1(\boldsymbol{\omega}_\ell).$$

"Power" method: Restrict A to the linear space

$$\operatorname{Span}(\mathbf{A}^{q}\boldsymbol{\omega}_{1}, \mathbf{A}^{q}\boldsymbol{\omega}_{2}, \ldots, \mathbf{A}^{q}\boldsymbol{\omega}_{\ell}).$$

Modified "power" method: Restrict A to the linear space

$$\mathcal{V}_q(oldsymbol{\omega}_1) imes \mathcal{V}_q(oldsymbol{\omega}_2) imes \cdots imes \mathcal{V}_q(oldsymbol{\omega}_\ell).$$

This could be a promising area for further work.

The observation that a "thin" Gaussian random matrix to high probability is well-conditioned is at the heart of the celebrated Johnson-Lindenstrauss lemma:

Lemma: Let ε be a real number such that $\varepsilon \in (0, 1)$, let *n* be a positive integer, and let *k* be an integer such that

(16)
$$k \ge 4 \left(\frac{\varepsilon^2}{2} - \frac{\varepsilon^3}{3}\right)^{-1} \log(n).$$

Then for any set V of n points in \mathbb{R}^d , there is a map $f : \mathbb{R}^d \to \mathbb{R}^k$ such that

(17)
$$(1-\varepsilon) ||u-v||^2 \le ||f(u)-f(v)|| \le (1+\varepsilon) ||u-v||^2, \quad \forall u, v \in V.$$

Further, such a map can be found in randomized polynomial time.

It has been shown that an excellent choice of the map f is the linear map whose coefficient matrix is a $k \times d$ matrix whose entries are i.i.d. Gaussian random variables (see, *e.g.* Dasgupta & Gupta (1999)). When k satisfies, (16), this map satisfies (17) with probability close to one.
The related Bourgain embedding theorem shows that such statements are not restricted to Euclidean space:

Theorem: Every finite metric space (X, d) can be embedded into ℓ^2 with distortion $O(\log n)$ where *n* is the number of points in the space.

Again, random projections can be used as the maps.

The Johnson-Lindenstrauss lemma (and to some extent the Bourgain embedding theorem) expresses a theme that is recurring across a number of research areas that have received much attention recently. These include:

- Compressed sensing (Candès, Tao, Romberg, Donoho).
- Approximate nearest neighbor search (Jones, Rokhlin).
- Geometry of point clouds in high dimensions (Coifman, Jones, Lafon, Lee, Maggioni, Nadler, Singer, Warner, Zucker, *etc*).
- Construction of multi-resolution SVDs.
- Clustering algorithms.
- Search algorithms / knowledge extraction.

Note: Omissions! No ordering. Missing references. Etc etc.

Many of these algorithms work "unreasonably well."

The randomized algorithm presented here is close in spirit to randomized algorithms such as:

• Randomized quick-sort.

(With variations: computing the median / order statistics / *etc*.)

- Routing of data in distributed computing with unknown network topology.
- Rabin-Karp string matching / verifying equality of strings.
- Verifying polynomial identities.

Many of these algorithms are of the type that it is the *running time* that is stochastic. The quality of the final output is excellent.

The randomized algorithm that is perhaps the best known within numerical analysis is Monte Carlo. This is somewhat lamentable given that MC is often a "last resort" type algorithm used when the curse of dimensionality hits — inaccurate results are tolerated simply because there are no alternatives.

(These comments apply to the traditional "unreformed" version of MC — for many applications, more accurate versions have been developed.)

Observation: Mathematicians working on these problems often focus on minimizing the distortion factor

$$\frac{1+\varepsilon}{1-\varepsilon}$$

arising in the Johnson-Lindenstrauss bound:

$$(1-\varepsilon) ||u-v||^2 \le ||f(u) - f(v)|| \le (1+\varepsilon) ||u-v||^2, \quad \forall u, v \in V.$$

In our environments, we do not need this constant to be particularly close to 1. It should just not be "large" — say less that 10 or some such.

This greatly reduces the number of random projections needed! Recall that in the Johnson-Lindenstrauss theorem:

number of samples required
$$\sim \frac{1}{\epsilon^2} \log(N)$$
.

Observation: Multiplication by a random unitary matrix reduces any matrix to its "general" form. All information about the singular vectors vanish. (The singular *values* remain the same.)

This opens up the possibility for general pre-conditioners — counterexamples to various algorithms can be disregarded.

The feasibility has been demonstrated for the case of least squares solvers for very large, very over determined systems. (Work by Rokhlin & Tygert, Sarlós,)

Work on $O(N^2 (\log N)^2)$ solvers of general linear systems is under way. (Random pre-conditioning + iterative solver.)

May stable fast matrix inversion schemes for general matrices be possible?

Observation: Robustness with respect to the quality of the random numbers.

The assumption that the entries of the random matrix are i.i.d. normalized Gaussians simplifies the analysis since this distribution is invariant under unitary maps.

In practice, however, one can use a low quality random number generator. The entries can be uniformly distributed on [-1, 1], they be drawn from certain Bernouilli-type distributions, *etc*.

Remarkably, they can even have enough internal structure to allow fast methods for matrix-vector multiplications. For instance:

- Subsampled discrete Fourier transform.
- Subsampled Walsh-Hadamard transform.
- Givens rotations by random angles acting on random indices.

This was exploited in "Algorithm 2" (and related work by Ailon and Chazelle). Our theoretical understanding of such problems is unsatisfactory. Numerical experiments perform *far* better than existing theory indicates.