
SketchySVD



Joel A. Tropp

Computing + Mathematical Sciences

California Institute of Technology

`jtropp@cms.caltech.edu`

Coauthors: Alp Yurtsever (EPFL), Volkan Cevher (EPFL);

Yiming Sun (Cornell), Yang Guo (UWisc), Charlene Luo (Columbia), Madeleine Udell (Cornell)

Thanks: Gunnar Martinsson (UT-Austin), Mark Tygert (Facebook)

Open for Business


**SIAM Journal on
Mathematics of Data Science
(SIMODS)**

<https://simods.siam.org/>

Editors: [Tammy Kolda \(EIC\)](#); Mark Girolami, Al Hero, Rob Nowak, Joel Tropp

The Famous Truncated SVD

Truncated Singular Value Decomposition (TSVD)



- U, V have orthonormal columns and Σ is nonnegative diagonal
- **Interpretation:** r -truncated SVD = optimal rank- r approximation
- Approximately $r(m + n)$ degrees of freedom

Applications:

- Least-squares computations (linear regression)
- Principal component analysis (orthogonal regression; total least squares)
- Summarization, data reduction, visualization, ...

A Pæan to the Truncated SVD

“[Truncated SVD] is one of the few methods that has solid foundations and can be trusted, provided that the computations are correct. Having something reliable in machine learning is worth its weight in gold — there are almost no gold standards in the field, precluding rapid progress. Think of what happens to machine learning when the routine for matrix–vector multiplication doesn’t always work right. You can’t debug any code, much less a big system consisting of many algorithms thrown together.”

—Nemo

Modern Numerical Linear Algebra

What's Wrong with Classical TSVD Algorithms?

- 🐛 Nothing... when the matrices are small and fit in core memory

Climate Change:

- 🐛 Medium- to large-scale data (Gigabytes+)
- 🐛 New architectures (multi-core, distributed, data centers, ...)
- 🐛 **Today:** New data presentations (off-core, dynamic, streaming)

The Role of Randomness:

- 🐛 Randomness is becoming a core tool in numerical linear algebra
- 🐛 Supports computations that are otherwise impossible
- 🐛 Numerical analysis remains essential!
- 🐛 **Today:** A practical randomized algorithm for streaming TSVD

History of Randomized SVD Algorithms

Dimension Reduction:

- 🐼 Gaussian maps (Johnson & Lindenstrauss 1984; Indyk & Motwani 1998)
- 🐼 **Sparse maps** (Achlioptas 2001; Charikar et al. 2002; Clarkson & Woodruff 2011; ...)
- 🐼 Randomized Fourier transforms (Ailon & Chazelle 2006; Woolfe et al. 2007)
- 🐼 **Tensor random projections** (Rudelson 2011; Sun et al. 2018)




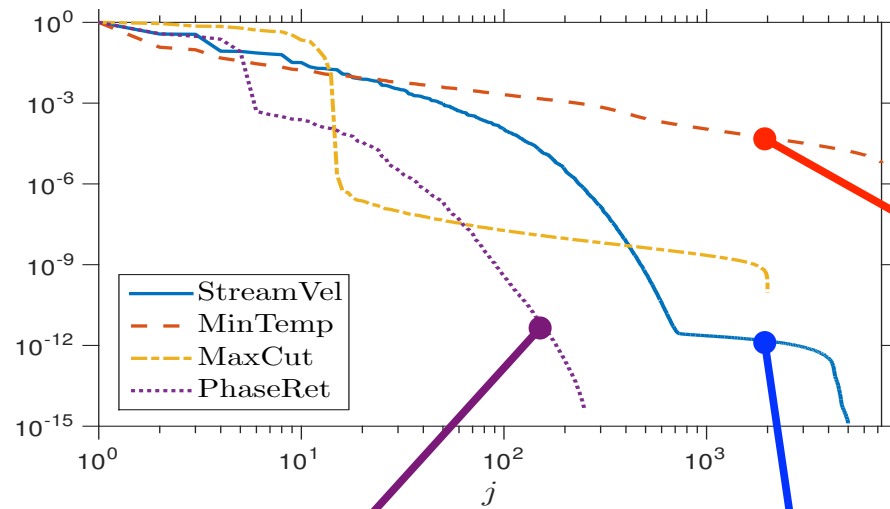
Theory:

- 🐼 Randomized algorithms for LSI (Frieze, Kannan, Papadimitriou, Vempala 1998)
- 🐼 Randomized linear algebra foundations (Drineas, Kannan, Mahoney 2004)
- 🐼 Sketch-and-solve framework (Sarlós 2006)
- 🐼 Streaming linear algebra foundations (Clarkson & Woodruff 2009)
- 🐼 **Streaming SVD from three sketches** (Upadhyay 2016)

Practice:

- 🐼 Low-rank matrix approximation algorithms (Martinsson, Rokhlin, Tygert 2004)
- 🐼 Fast, one-pass matrix approximation (Woolfe, Liberty, Rokhlin, Tygert 2007)
- 🐼 **Randomized SVD framework, algorithms, and analysis** (Halko, Martinsson, Tropp 2009)
- 🐼 Randomized block Krylov methods (Halko, Martinsson, Shkolnisky, Tygert 2011)
- 🐼 **Practical streaming SVD algorithms** (Tropp, Yurtsever, Udell, Cevher 2016–2019)

Spectral Decay in Scientific Data



Streaming Linear Algebra

The Turnstile Model:

$$\mathbf{A} = \mathbf{H}_1 + \mathbf{H}_2 + \mathbf{H}_3 + \mathbf{H}_4 + \dots$$

- ☞ Huge input matrix \mathbf{A} is presented as a sum of innovations \mathbf{H}_i
- ☞ Must discard each innovation \mathbf{H}_i after it is processed
- ☞ **Goal:** Without storing \mathbf{A} in full, return TSVD after seeing all updates

Applications:

- ☞ One-pass approximation of matrix stored out-of core
- ☞ Large-scale semidefinite programming algorithms
- ☞ Scientific simulation and data collection

Sources: Muthukrishnan 2008; Woolfe et al. 2008; Clarkson & Woodruff 2009; HMT 2011; Woodruff 2014; TYUC 2016–2019;

Randomized Linear Sketches

$$\text{sketch} = \mathcal{L}(A) = \sum_i \mathcal{L}(H_i)$$

- ☞ Select a **linear** map \mathcal{L} without reference to A
- ☞ Sketch dimension is much smaller than input matrix dimension
- ☞ Use **randomness** so sketch works for an arbitrary input
- ☞ **[LNW14] Essentially the only way to handle the turnstile model!**


Examples:

- ☞ **Left multiply:** $\mathcal{L}(A) = \Upsilon A$ where Υ is fat
- ☞ **Right multiply:** $\mathcal{L}(A) = A\Omega$ where Ω is tall
- ☞ **Compress:** $\mathcal{L}(A) = \Phi A\Psi$ where Φ is fat, Ψ is tall

Sources: Alon et al. 1996; Sarlós 2006; Muthukrishnan 2008; Woolfe et al. 2008; Clarkson & Woodruff 2009; HMT 2011; Mahoney 2012; Woodruff 2014; Li et al. 2014; Drineas & Mahoney 2016; TYUC 2016–2019;

The Randomized Range Finder

Images of Random Vectors Align with the Range



The Randomized Range Finder

Q. How do we use a randomized linear sketch to approximate the range of A ?

A. Just multiply k random vectors into A and orthogonalize!

1. Draw a random matrix Ω with k columns
2. Form the range sketch: $Y = A\Omega$
3. Orthogonalize the columns of the sketch: $Q = \text{orth}(Y)$

Claim: $A \approx QQ^*A$

Slogan: “range(Q) is a k -dimensional subspace that captures the action of A ”

Sources: NLA community: Stewart (1970s). GFA: Johnson & Lindenstrauss (1984) et seq. TCS: Boutsidis, Cohen, Deshpande, Drineas, Frieze, Kannan, Mahoney, Nelson, Papadimitriou, Sarlós, Vempala (1998–present). SciComp: Halko, Liberty, Martinsson, Rokhlin, Szlam, Tropp, Tygert, Woolfe (2004–present). See HMT 2011, §2 for more history.

Analysis of the Randomized Range Finder

Theorem 1 (HMT 2011). Assume

- The input matrix $A \in \mathbb{C}^{m \times n}$
- Range sketch $Y = A\Omega$ where $\Omega \in \mathbb{C}^{n \times k}$ is *complex standard normal*
- Orthogonalize columns: $Q = \text{orth}(Y)$

Then

$$\mathbb{E} \left\| A - QQ^* A \right\|_F \leq \sqrt{2} \cdot \left\| A - [A]_{k/2} \right\|_F$$

- **Key Fact:** Approximation exploits spectral decay
- There's nothing special about the choice $k/2$
- Probability of a much larger error is negligible
- Related results hold for the spectral norm

Source: Halko et al. 2011, §10.2; Gu 2015; TYUC 2016–2019.

Overview of SketchySVD

A Tripartite Sketch

☞ Let $A \in \mathbb{C}^{m \times n}$ be an input matrix (presented in turnstile model)

☞ Fix sketch size parameter k with $r \leq k \ll \min\{m, n\}$

☞ Draw independent random linear maps:

$$\begin{aligned} \Upsilon &\in \mathbb{C}^{k \times m} & \text{and} & & \Omega &\in \mathbb{C}^{n \times k} \\ \Phi &\in \mathbb{C}^{2k \times m} & \text{and} & & \Psi &\in \mathbb{C}^{n \times 2k} \end{aligned}$$

☞ **Sketches for right and left dominant singular subspaces:**

$$X = \Upsilon A \in \mathbb{C}^{k \times n} \quad \text{and} \quad Y = A \Omega \in \mathbb{C}^{m \times k}$$

☞ **Sketch to estimate the core:**

$$Z = \Phi A \Psi \in \mathbb{C}^{2k \times 2k}$$

Sources: Vempala et al. 1998–2000; Drineas et al. 2004–2006; Martinsson et al. 2004–2012; Sarlós 2006; Woolfe et al. 2008; Clarkson & Woodruff 2009, 2011; Boutsidis et al. 2011, 2016; HMT 2011; Mahoney 2012; Nelson et al. 2012–2015; Woodruff 2014; Gu 2015; Cohen et al. 2015; Upadhyay 2016; TYUC 2016–2019....

Pseudocode for SKETCHYSVD

Input: Sketch size parameter k ; input matrix $A \in \mathbb{C}^{m \times n}$ as a turnstile stream

Output: Rank- r approximation $\hat{A}_r = \mathbf{U}\Sigma\mathbf{V}^*$

```
1  function INITIALIZE( $m, n, k$ )                                ▷ Set up the sketch
2      Draw random linear maps  $\Upsilon, \Omega, \Phi, \Psi$ 
3       $X \leftarrow \mathbf{0}$  and  $Y \leftarrow \mathbf{0}$  and  $Z \leftarrow \mathbf{0}$ 
4  function LINEARUPDATE( $H$ )                                    ▷ Process  $A \leftarrow A + H$ 
5       $X \leftarrow X + \Upsilon H$ 
6       $Y \leftarrow Y + H\Omega$ 
7       $Z \leftarrow Z + \Phi H\Psi$ 
8  function SKETCHYSVD( $r$ )                                       ▷ Compute  $r$ -truncated SVD
9       $Q \leftarrow \text{thin\_qr}(Y)$                                 ▷ Basis for range
10      $P \leftarrow \text{thin\_qr}(X^*)$                                ▷ Basis for co-range
11      $C \leftarrow ((\Phi Q) \setminus Z) / (P^* \Psi)$              ▷ Core matrix
12      $(U, \Sigma, V) \leftarrow \text{svd}(C; r)$                    ▷ Truncate dense SVD of core
13      $U \leftarrow QU$  and  $V \leftarrow PV$                    ▷ Consolidate unitary factors
14     return  $(U, \Sigma, V)$ 
```

Resource Usage with Sparse Maps

Storage:

- 🐼 Random linear maps: $O(m + n)$
- 🐼 Sketches: $O(k(m + n))$

Arithmetic:

- 🐼 Linear update: Depends on structure of update (cheap!)
- 🐼 SKETCHYSVD: $O(k^2(m + n))$

Sources: Charikar et al. 2002; Cormode & Muthukrishnan 2005; Sarlós 2006; Woolfe et al. 2008; Clarkson & Woodruff 2009, 2011; HMT 2011; Nelson et al. 2012–2015; Meng & Mahoney 2013; Cohen 2015; TYUC 2016–2019....

Analysis of SKETCHYSVD

Theorem 2 (TYUC 2018). **Assume**

• The input matrix $\mathbf{A} \in \mathbb{C}^{m \times n}$

• The dimension reduction maps are *independent complex standard normal*

Then SKETCHYSVD computes a rank- r approximation $\hat{\mathbf{A}}_r$ for which

$$\mathbb{E} \|\mathbf{A} - \hat{\mathbf{A}}_r\|_{\text{F}} \leq \|\mathbf{A} - [\mathbf{A}]_r\|_{\text{F}} + 5 \cdot \|\mathbf{A} - [\mathbf{A}]_{k/2}\|_{\text{F}}$$

• **Key Fact:** Approximation exploits spectral decay

• Probability of a much larger error is negligible

• Related results hold for the spectral norm

Source: HMT 2011; Gu 2015; TYUC 2016–2019.


Empirical Performance

Important Things I'm Not Going to Show You

- 🦉 SKETCHYSVD is insensitive to the distributions of the random linear maps
- 🦉 Theory gives parameter choices that are nearly optimal in practice
- 🦉 SKETCHYSVD beats earlier techniques for synthetic and real data
- 🦉 Why we need to truncate the approximation
- 🦉 Methodology for estimating errors and selecting the truncation rank
- 🦉 Sampling distribution of approximation error and error estimator
- 🦉 Other structured approximations via re-factorization or matrix nearness
- 🦉 Transformation of the data before sketching
- 🦉 Extension to low-rank Tucker approximation of a tensor

Sources: HMT 2011; TYUC 2016–2019; SGLTU 2018–2019.


Reconstruction of von Kármán Street




Comments: Data: $m = 10,738$; $n = 5,001$; 430 MB. Algorithm: [TYUC19], sparse maps; rank $r = 5$; storage $T = 48(m + n)$. Compression: $71 \times$.

Left Singular Vectors of von Kármán Street

Approximate [TYUC19]





Exact





Comments: Data: $m = 10,738$; $n = 5,001$; 430 MB. Algorithm: Sparse maps; rank $r = 5$; storage $T = 48(m + n)$. Compression: $71\times$.



Singular Vectors of Sea Surface Temperature Data





Spatiotemporal Avg.





Seasonal



(Intra-)Seasonal



(Intra-)Seasonal



El Niño / La Niña

Comments: Data: $m = 691,150$; $n = 13,670$; 75 GB. Algorithm: Sparse maps; $k = 48$; $s = 839$. Compression ratio: $222\times$.

To learn more...

E-mail: jtropp@cms.caltech.edu

Web: <http://users.cms.caltech.edu/~jtropp>


Papers:

- Halko, Martinsson, & Tropp, “Finding structure with randomness: Probabilistic algorithms for computing approximate matrix decompositions,” *SIREV*, 2011.
- Tropp, Yurtsever, Udell, & Cevher, “Sketchy decisions: Low-rank matrix optimization with optimal storage,” AISTATS 2017.
- Tropp, Yurtsever, Udell, & Cevher, “Practical sketching algorithms for low-rank matrix approximation,” *SIMAX*, 2017.
- Tropp, Yurtsever, Udell, & Cevher, “Fixed-rank approximation of a positive-semidefinite matrix from streaming data,” *NeurIPS*, 2017.
- Tropp, Yurtsever, Udell, & Cevher, “Streaming low-rank matrix approximation with an application to scientific simulation,” arXiv cs.NA 1902.08651. *SISC*, 2019.
- Sun, Guo, Tropp, & Udell, “Tensor random projections for low-memory dimension reduction,” *NeurIPS Relational Databases Workshop*, 2018.
- Sun, Guo, Luo, Tropp, & Udell, “Low-rank Tucker approximation of a tensor from streaming data,” arXiv cs.NA 1904.10951. *SIMODS*, 2020.
- Cevher, Tropp, & Yurtsever, “Scalable semidefinite programming,” arXiv math.OC 1912.02949, 2019.


Supplementary Materials

Insensitivity to Dimension Reduction Map


LowRankHiNoise
 LowRankMedNoise
 LowRankLowNoise



PolyDecaySlow
 PolyDecayMed
 PolyDecayFast




ExpDecaySlow
 ExpDecayMed
 ExpDecayFast




Comments: Effective rank $R = 10$, approximation rank $r = 10$, Schatten 2-norm.

Performance with Theoretical Parameter Choices


LowRankHiNoise
 LowRankMedNoise
 LowRankLowNoise



PolyDecaySlow
 PolyDecayMed
 PolyDecayFast






ExpDecaySlow
 ExpDecayMed
 ExpDecayFast






Comments: Gaussian maps, effective rank $R = 10$, approximation rank $r = 10$, Schatten 2-norm.

Method Comparison: Synthetic Data




LowRankHiNoise
 LowRankMedNoise
 LowRankLowNoise



PolyDecaySlow
 PolyDecayMed
 PolyDecayFast




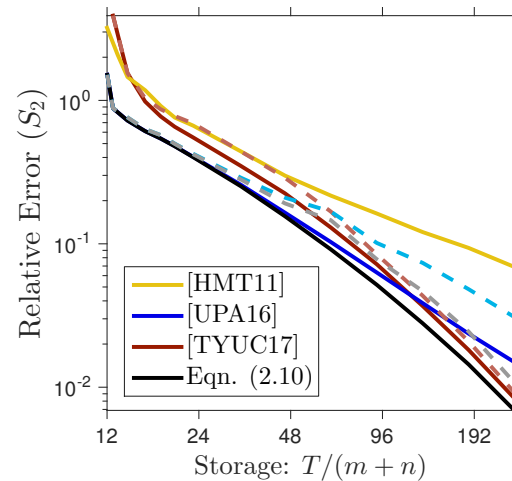
ExpDecaySlow
 ExpDecayMed
 ExpDecayFast



Comments: Gaussian maps, effective rank $R = 10$, oracle parameters, approximation rank $r = 10$, Schatten 2-norm.)



Method Comparison: Real Data

MinTemp
 $m = 19,264$
 $n = 7,305$
 $r = 10$



StreamVel
 $m = 10,738$
 $n = 5,001$
 $r = 10$


MaxCut
 $n = 2,000$
 $r = 1$




PhaseRetrieval
 $n = 25,921$
 $r = 5$

Comments: Sparse maps, Schatten 2-norm. Solid lines are errors with oracle parameters; dashed lines are *a priori* parameter choices.

Why Truncate?



(A) Spectrum of Approximation $\sigma_r(\hat{A}_k)$



(B) Error in Rank- r Truncation \hat{A}_r

$$\text{relerr}(\hat{A}_r) = \frac{\|\mathbf{A} - \hat{A}_r\|}{\|\mathbf{A} - \llbracket \mathbf{A} \rrbracket_r\|} - 1$$

Comments: StreamVel Data: $m = 10,738$; $n = 5,001$; 430 MB. Algorithm: Sparse maps; $s = 2k + 1$.

A Posteriori Error Estimation


- Fix a sketch size parameter q
- Draw a random Gaussian dimension reduction map $\Theta \in \mathbb{C}^{m \times q}$
- Maintain an error sketch $\mathbf{S} = \Theta \mathbf{A}$
- Given an approximation $\hat{\mathbf{A}}$, compute the error estimator

$$\text{err}_2^2(\hat{\mathbf{A}}) = \|\mathbf{S} - \Theta \hat{\mathbf{A}}\|_F^2$$


- The error estimator is unbiased and concentrates sharply
- We can also compute an empirical upper bound on the scree curve as

$$\overline{\text{scree}}(r) = \left[\frac{\tau_{r+1}(\hat{\mathbf{A}}) + \text{err}_2(\hat{\mathbf{A}})}{\text{err}_2(\mathbf{0})} \right]^2.$$


Error Estimates and Empirical Scree Curves




(A) Error Estimates for \hat{A}



(B) Scree Plot ($k = 16$)




(C) Scree Plot ($k = 48$)




(D) Scree Plot ($k = 128$)

Comments: StreamVel, sparse maps, $s = 2k + 1$, $q = 10$.


Sampling Distribution of Error and Estimator




(A) Rank- k Approximation ($k = 48$)



(B) Rank- r Truncation ($k = 48, r = 12$)




(C) Rank- k Approximation ($k = 128$)



(D) Rank- r Truncation ($k = 128, r = 32$)

Comments: StreamVel, sparse maps, $s = 2k + 1$.

Scree Plot for Sea Surface Temperature Data



Comments: SeaSurfaceTemp, sparse maps, $k = 48$, $s = 839$, $q = 10$