Faster Linear Algebra Algorithms with Structured Random Matrices

Chris Camaño, Ethan N. Epperly, Raphael A. Meyer, Joel A. Tropp

Caltech

Introduction

For many years, the randomized numerical linear algebra community has increasingly favored structured sketching matrices, such as sparse embeddings and randomized trigonometric transforms over their (dense) Gaussian counterparts.

Despite their empirical success, these matrices are often used without strong theoretical guarantees. In this work, we relax the standard oblivious subspace embedding (OSE) condition and demonstrate that many structured random matrices can be shown to produce near-optimal error using only bounds on the minimal singular value. With these new tools we confirm a long held suspicion: **Structured test matrices provide the same accuracy as Gaussian test matrices, while the structure allows us to design far more efficient algorithms.**

Oblivious Subspace Embedding (OSE)

A random matrix $\mathbf{\Omega} \in \mathbb{F}^{d\times k}$ is called an (r,α,β)-OSE with subspace dimension r, embedding dimension $k \ge r$, injectivity $\alpha \in (0,1]$, and dilation $\beta \ge 1$ if the following holds for each fixed r-dimensional subspace $\mathscr{V} \subseteq \mathbb{F}^d$. With probability at least 19/20,

$$\alpha\|\boldsymbol{x}\|_2^2 \le \|\mathbf{\Omega}^*\boldsymbol{x}\|_2^2 \le \beta\|\boldsymbol{x}\|_2^2 \quad \text{for all } \boldsymbol{x} \in \mathscr{V}.$$

Typically, a sketching-based randomized algorithm begins by assuming that the random matrix satisfies the OSE property.

Oblivious Subspace *Injection* (OSI)

A random matrix $\mathbf{\Omega} \in \mathbb{F}^{d\times k}$ is called an (r,α)-OSI with subspace dimension r, embedding dimension $k \ge r$, and injectivity $\alpha \in (0,1]$ when it meets two conditions:

1. Isotropy. On average, the matrix preserves the squared length of each vector:

$$\mathbb{E}\left\|\mathbf{\Omega}^*\boldsymbol{x}\right\|_2^2 = \|\boldsymbol{x}\|_2^2 \quad \text{for all } \boldsymbol{x} \in \mathbb{F}^d.$$

2. Injectivity. For each fixed r-dimensional subspace $\mathscr{V} \subseteq \mathbb{F}^d$, with probability at least $\frac{19}{20}$,

$$\alpha \cdot \|\boldsymbol{x}\|_2^2 \le \|\mathbf{\Omega}^*\boldsymbol{x}\|_2^2 \quad \text{for all } \boldsymbol{x} \in \mathscr{V}.$$

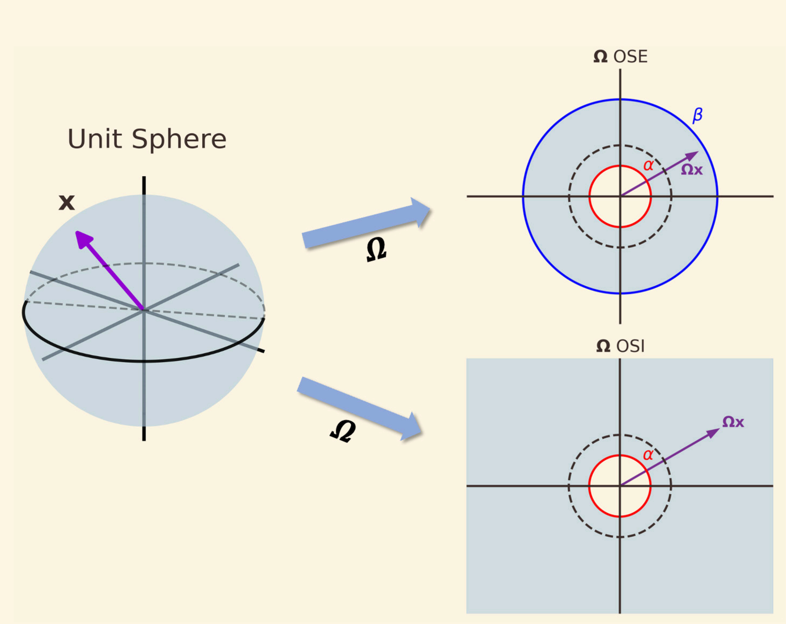

Intuition: Dilation versus injection. Let $\boldsymbol{\omega} \in \mathbb{F}^d$ be isotropic with $\mathbb{E}[\boldsymbol{\omega}\boldsymbol{\omega}^*] = \mathbf{I}$, and define

$$\mathbf{\Omega} := \frac{1}{\sqrt{k}}[\,\boldsymbol{\omega}_1 \ \cdots \ \boldsymbol{\omega}_k\,] \in \mathbb{F}^{d\times k}, \qquad \boldsymbol{\omega}_j \overset{\text{iid}}{\sim} \boldsymbol{\omega}.$$

For any fixed $\boldsymbol{x} \in \mathbb{F}^d$,

$$\|\mathbf{\Omega}^*\boldsymbol{x}\|_2^2 = \frac{1}{k}\sum_{i=1}^{k}\left|\langle\boldsymbol{\omega}_i,\boldsymbol{x}\rangle\right|^2,$$

is an average of iid nonnegative terms with mean $\|\boldsymbol{x}\|_2^2$. Dilation is driven by a single unusually large summand (heavy-tailed marginals inflate β), injection failure requires all summands to be annihilated.

Spectral formulation. An equivalent statement of the OSI condition can be expressed in spectral terms. For any fixed matrix $\boldsymbol{Q} \in \mathbb{F}^{d\times r}$ with orthonormal columns, suppose

$$\sigma_{\min}^2(\mathbf{\Omega}^*\mathbf{Q}) = \lambda_{\min}(\mathbf{Q}^*\mathbf{\Omega}\mathbf{\Omega}^*\mathbf{Q}) \ge \alpha > 0 \quad \text{with probability at least 19/20.}$$

Then $\mathbf{\Omega}$ satisfies the OSI condition.

Structured Random Matrices Are Just as Good as Gaussians

SparseStacks

The **SparseStack** test matrix is determined by the ambient dimension d, row sparsity ζ, and block size b, giving embedding dimension $k = b\zeta$. It is defined blockwise as

$$\mathbf{\Omega} := \frac{1}{\sqrt{\zeta}}\begin{bmatrix} \varrho_{11}\mathbf{e}_{s_{11}}^* & \cdots & \varrho_{1\zeta}\mathbf{e}_{s_{1\zeta}}^* \\ \varrho_{21}\mathbf{e}_{s_{21}}^* & \cdots & \varrho_{2\zeta}\mathbf{e}_{s_{2\zeta}}^* \\ \vdots & \ddots & \vdots \\ \varrho_{d1}\mathbf{e}_{s_{d1}}^* & \cdots & \varrho_{d\zeta}\mathbf{e}_{s_{d\zeta}}^* \end{bmatrix} \in \mathbb{F}^{d\times k}, \quad \begin{array}{l}\varrho_{ij} \sim \text{Rademacher}, \\ s_{ij} \sim \text{Uniform}\{1,\dots,b\}.\end{array}$$

Here $\mathbf{e}_i \in \mathbb{F}^b$ is the ith standard basis vector. A **SparseStack** is equivilant to the row tiling of ζ CountSketches.

- $(r, 1/2)$-OSI guarantee with $k = \mathcal{O}(r)$ and $\zeta = \log(r)$ row sparsity
- **Sketching time**: For $\mathbf{A} \in \mathbb{F}^{n\times d}$, the sketch $\mathbf{A\Omega}$ requires only $\mathcal{O}(\zeta \operatorname{nnz}(\mathbf{A}))$ operations

Results for Randomized SVD

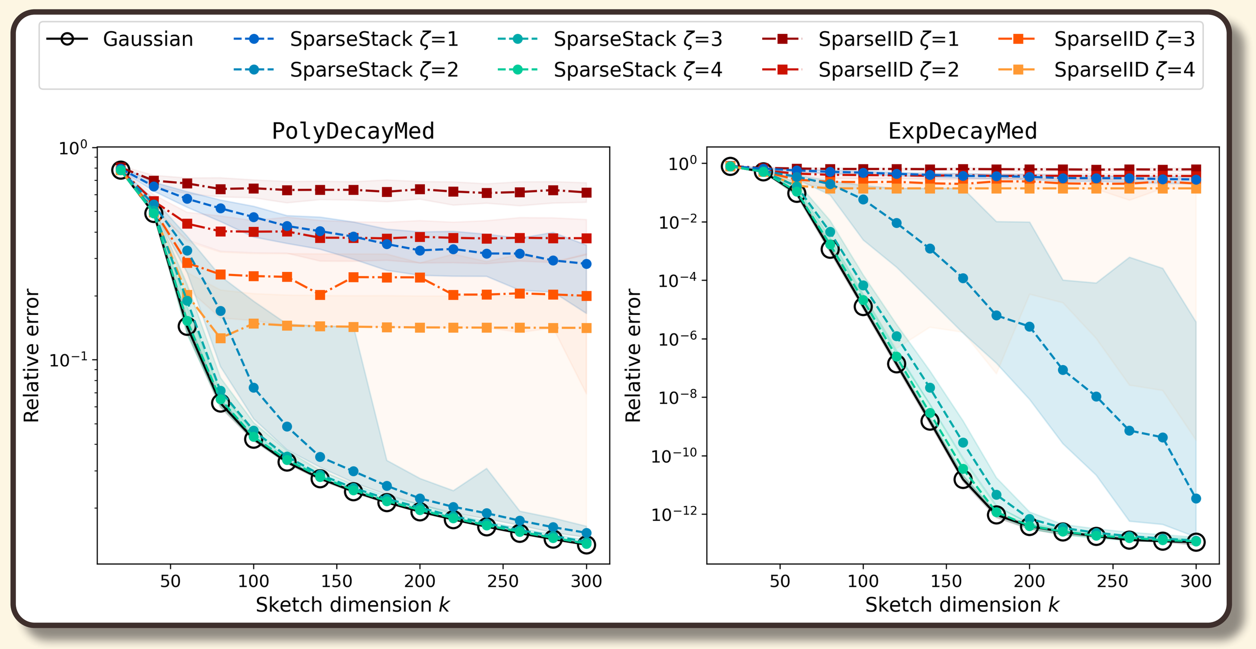

Sparse Randomized Trigonometric Transform (SparseRTT)

The **SparseRTT** test matrix is defined as

$$\mathbf{\Omega} := \boldsymbol{DFS} \in \mathbb{F}^{d\times k} \quad \text{where} \quad \begin{array}{l} \boldsymbol{D} \in \mathbb{F}^{d\times d} \text{ is a random diagonal matrix,} \\ \boldsymbol{F} \in \mathbb{F}^{d\times d} \text{ is a trigonometric transform,} \\ \boldsymbol{S} \in \mathbb{F}^{d\times k} \text{ is a SparseCol matrix.} \end{array}$$

A **SparseCol** a sparse random matrix with ξ nnz enteries per column chosen without replacement.

- $(r, 1/2)$-OSI guarantee with $k = \mathcal{O}(r)$ and column sparsity $\zeta = \mathcal{O}(\log r)$
- **Sketching time**: For $\mathbf{A} \in \mathbb{F}^{n\times d}$, the sketch $\mathbf{A\Omega}$ requires $\mathcal{O}(nd \log r)$ operations

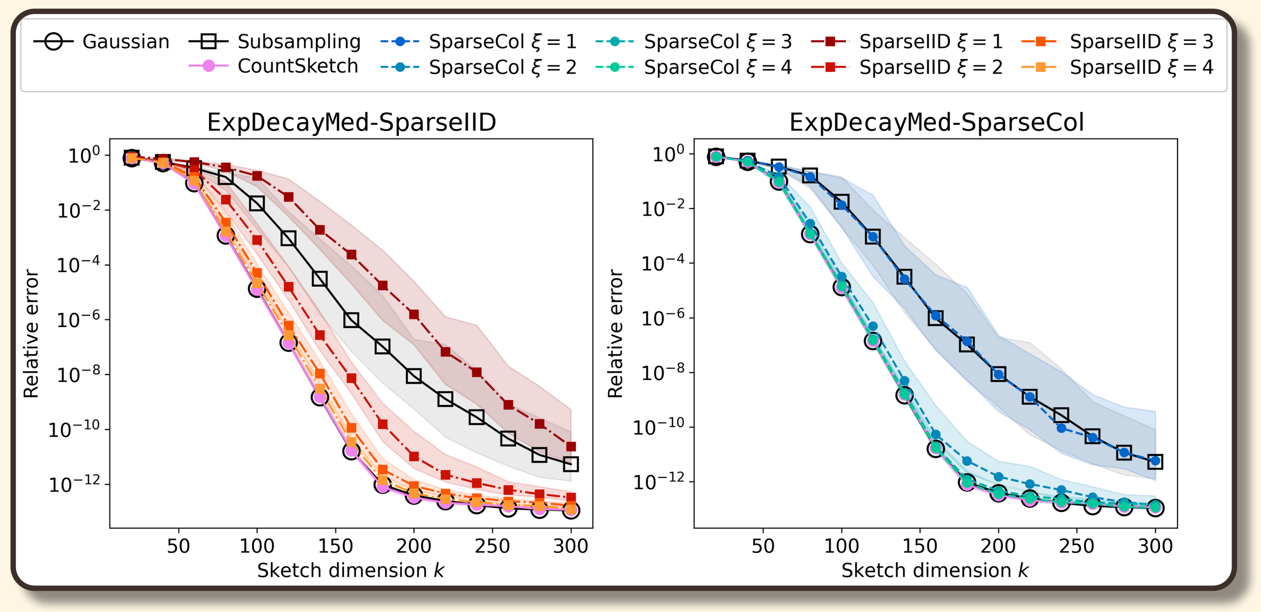

Khatri-Rao

The **Khatri–Rao** test matrix is constructed from an isotropic base distribution $\boldsymbol{v} \in \mathbb{F}^{d_0}$. Let ℓ denote the Kronecker depth and k the embedding dimension. Then the Khatri–Rao test matrix is defined as

$$\mathbf{\Omega} = \frac{1}{\sqrt{k}}[\,\boldsymbol{\omega}_1 \ \cdots \ \boldsymbol{\omega}_k\,] \in \mathbb{F}^{d_0^\ell \times k}, \qquad \boldsymbol{\omega}_i := \boldsymbol{\omega}_i^{(1)} \otimes \cdots \otimes \boldsymbol{\omega}_i^{(\ell)}, \qquad \boldsymbol{\omega}_i^{(j)} \overset{\text{iid}}{\sim} \boldsymbol{v}.$$

- $(r, 1/2)$-OSI guarantee with $k = \mathcal{O}(C_v^\ell r)$ for Kronecker depth ℓ where C_v is a distribution dependent constant.
- **Sketching time**: In the Kronecker matvec access model, $\mathbf{A} \in \mathbb{R}^{d^\ell \times d^\ell}$ can be sketched efficiently using tensor networks in $\mathcal{O}(n\,k\,\ell^2)$ time, avoiding dense Gaussian embeddings.

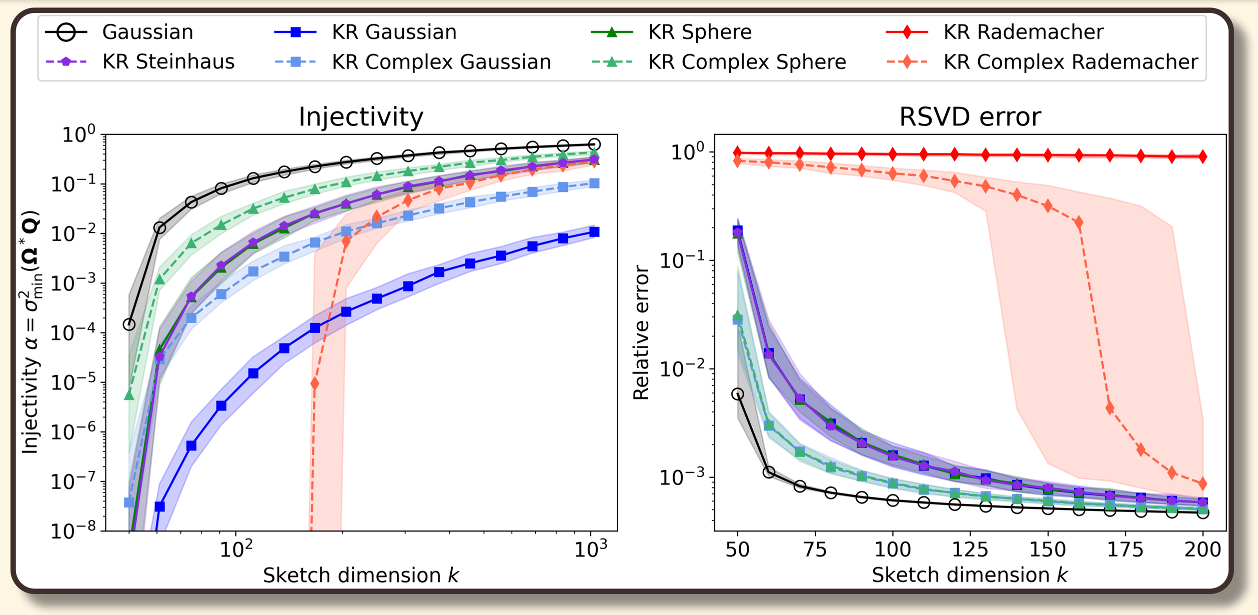

Improvements Over Other Sparse Sketches

Table 1: **Sparse test matrices: Comparison.** Theoretical results for sparse test matrices, including the type of guarantee (OSE or OSI), the embedding dimension k and row sparsity ζ for sketching an r-dimensional subspace, and the runtime for generalized Nyström .

Guarantee	Test matrix	Embedding dim. k	Row sparsity ζ	Gen. Nyström runtime
OSE	CountSketch [Nelson et al. 13']	$\mathcal{O}(r^2)$	1	$\mathcal{O}(\operatorname{nnz}(\mathbf{A}) + nr^4)$
	SparseStack [Cohen 16']	$\mathcal{O}(r\log r)$	$\mathcal{O}(\log r)$	$\mathcal{O}(\operatorname{nnz}(\mathbf{A})\log(r) + nr^2\log^3 r)$
	SparseStack [Chenakkod et al. 24']	$\mathcal{O}(r)$	$\mathcal{O}(\log^3 r)$	$\mathcal{O}(\operatorname{nnz}(\mathbf{A})\log^3(r) + nr^2)$
OSI	SparseIID [Tropp 25']	$\mathcal{O}(r)$	$\mathcal{O}(\log r)$	$\mathcal{O}(\operatorname{nnz}(\mathbf{A})\log(r) + nr^2)$
	SparseStack (This work)	$\mathcal{O}(r)$	$\mathcal{O}(\log r)$	$\mathcal{O}(\operatorname{nnz}(\mathbf{A})\log(r) + nr^2)$
None	SparseStack (Practice)	$2r$	4	$\mathcal{O}(\operatorname{nnz}(\mathbf{A}) + nr^2)$

References

- T. Sarlos, *Improved Approximation Algorithms for Large Matrices via Random Projections*, FOCS'06, IEEE, 2006.
- J. Nelson and H. L. Nguyễn, *OSNAP: Faster Numerical Linear Algebra Algorithms via Sparser Subspace Embeddings*, FOCS'13, IEEE, 2013.
- J. A. Tropp, *ACM 217: Probability in High Dimensions*, CMS Lecture Notes, Caltech, 2021.
- J. Chenakkod, M. Dereziński, X. Dong, and M. Rudelson, *Optimal Embedding Dimension for Sparse Subspace Embeddings*, STOC'24, ACM, 2024.
- (This work) **C. Camaño, E. N. Epperly, R. A. Meyer, and J. A. Tropp, *Faster Linear Algebra Algorithms with Structured Random Matrices*, arXiv, 2025.**

Supported by the Caltech , the ONR, NSF GRFP, and DOE CSGF

Arxiv Link to paper