BIT Numerical Mathematics (2026) 66:7
https://doi.org/10.1007/s10543-025-01101-9

BIT

$S^\top S$-SVD via sketching and the nearest $S^\top S$-orthogonal matrix

Davide Palitta[1,2] · **Valeria Simoncini**[1,2]

Received: 30 May 2025 / Accepted: 6 December 2025

Abstract

Sketching techniques have gained popularity in numerical linear algebra to accelerate the solution of least squares problems. The so-called ε-subspace embedding property of a sketching matrix S has been largely used to characterize the problem residual norm, since the procedure is no longer optimal in terms of the (classical) Frobenius or Euclidean norm. By building on available results on the SVD of the sketched matrix SA derived by Gilbert, Park, and Wakin (Proc. of SPARS-2013), a novel decomposition of A, the $S^\top S$-SVD, is proposed, which *holds* with high probability, and in which the left singular vectors are orthonormal with respect to a (semi-)norm defined by the sketching matrix S. The new decomposition is less expensive to compute than the standard SVD, while preserving the singular values with probabilistic confidence. The $S^\top S$-SVD appears to be the right tool to analyze the quality of several sketching-based techniques in the literature, for which examples are reported. For instance, it is possible to simply bound the distance from (standard) orthogonality of sketching-based orthogonal matrices in state-of-the-art randomized algorithms for QR factorizations. As an application, the classical problem of the nearest orthogonal matrix is generalized to the new $S^\top S$-orthogonality, and the $S^\top S$-SVD is used to solve it. Probabilistic bounds on the quality of the solution are also derived.

Keywords Sketching · SVD · Nearest orthogonal matrix

Mathematics Subject Classification 65F99 · 68W20

This paper is dedicated to Åke Björck, on the occasion of his 90th birthday.

✉ Valeria Simoncini
valeria.simoncini@unibo.it

Davide Palitta
davide.palitta@unibo.it

[1] Dipartimento di Matematica and (AM)^2, Alma Mater Studiorum Università di Bologna, Piazza di Porta San Donato 5, I-40127 Bologna, Italy

[2] IMATI-CNR, Pavia, Italy

Published online: 05 January 2026

1 Introduction

The inspiration for this work comes from [7] where, for a given matrix $A \in \mathbb{R}^{m\times n}$, Björck and Bowie studied the problem

$$\min_{\substack{Q\in\mathbb{R}^{m\times n}\\ Q^\top Q=I}} \|A-Q\|_*, \tag{1}$$

with $m \geq n$ and $* = 2, F$, that is, the spectral and Frobenius norms. The solution Q to (1) is given by the orthogonal factor of the polar decomposition ([6, Th.1.2.4]) of A; see [7], and, e.g., the presentation and bibliography in [22]. However, this decomposition may be expensive to compute for large-scale problems, hence Björck and Bowie proposed an iterative algorithm to solve (1); see [7, Sect. 2].

With the same desire of accelerating the solution of (1) when A has large rank, we instead explore the use of state-of-the-art randomization-based strategies known as sketchings. Since the seminal work [28], sketching techniques have shown their potential in speeding up the numerical solution of massively overdetermined least squares problems; see, e.g., [2, 26]. These strategies reduce the row dimension of the coefficient matrix by projection, using suitable linear maps. Under certain probabilistic conditions, these maps can be ensured to have good metric properties. More precisely, for $\varepsilon \in (0, 1)$ and for any vector v in a k-dimensional vector space $\mathcal{V} \subset \mathbb{R}^m$, it is possible to construct a linear map, called an *oblivious ε-subspace embedding* $S : \mathbb{R}^m \mapsto \mathbb{R}^s$, $s \ll m$, such that $|\|Sv\|^2 - \|v\|^2| \leq \varepsilon$ with high probability. We remark that S can be selected without knowing $\mathcal{V}$ itself but relying only on its dimension k. Applying sketching to a least squares problem can be interpreted as recasting the original problem in terms of a different norm, the $S^\top S$-norm $\|v\|^2_{S^\top S} = v^\top S^\top S v$. Although $S^\top S$ is only semidefinite in general, the above metric property ensures that with high probability $S^\top S$ defines a positive definite norm on the embedded space; see, e.g., [4].

We aim to explore this non-standard norm and its theoretical and computational properties in the solution of (1). To this end, by building upon available results on the SVD of SA [13], we propose a novel decomposition of A that we name $S^\top S$-SVD, which *holds* with high probability. This amounts to an SVD-like decomposition of A, where the left singular vectors w_js are $S^\top S$-orthonormal, namely $w_j^\top S^\top S w_i = \delta_{ij}$, $i, j = 1, \ldots, n$; see Sect. 3. The magnitude of the $S^\top S$-singular values of A can be related to their standard counterparts via the threshold ε associated with the adopted embedding S [13]. This is one of the key features of the $S^\top S$-SVD that allow us to analyze the sketched version of (1). Moreover, we can derive explicit relations between the solution to the problem in the non-standard norm and the solution Q to (1).

The idea of sketching a matrix's SVD is certainly not new, and it has been introduced together with the development of sketching algorithms [18, 25, 28]. However, to our knowledge, *sketched SVD* usually refers to the construction of a reduced surrogate of A, built to preserve with high probability some of the properties of the singular values of A and of the range of $A^\top$ [13, 26]. Alternatives include randomized SVD associated with range-finding strategies; see [24] and the references therein. This work provides a new point of view on the use of SVD-related sketching techniques, as it aims at constructing a full decomposition of A, by computing an $S^\top S$-orthogonal basis for the range space. Indeed, the $S^\top S$-SVD computes a factorization of A, holding with high probability. By relying on the concept of $S^T S$-orthogonality, we are able to leverage the computational gains of sketching while attaining full factorizations with probabilistic confidence.

We believe that the $S^\top S$-SVD will be a crucial tool also in attaining a complete understanding of sketching techniques applied to general least squares problems. For instance, several results available in the literature are natural consequences of the $S^\top S$-SVD. To evaluate the quality of the new orthogonality constraint, we also estimate the distance from Euclidean orthogonality of the computed $S^\top S$-orthogonal matrices. These can be used, for instance, to provide certain orthogonality guarantees in state-of-the-art randomized algorithms for QR factorizations.

As possible motivating applications, we envision the employment of the $S^\top S$-SVD also in other settings. As an example, orthogonal matrices are a fundamental ingredient in problems where the solution is constrained to belong to a space of matrices having orthonormal columns, the Stiefel manifold; see, e.g., [1, 12]. These types of constraints are particularly convenient in the solution of certain differential equations, because the orthogonal space allows the solution method to preserve key properties of the dynamical system [10, 17]. Beyond this, orthogonality may be a component of computational strategies enforcing a low-rank manifold representation of the flow as time integration proceeds, see, e.g., [11]. The use of sketched orthogonality, that is of a “sketched” Stiefel manifold, may help decrease the computational costs of these procedures when large systems of differential equations arise. In other words, this strategy may be viewed as an ε-controlled quasi-orthogonality structure, and may lead to the supervised relaxation of orthogonal-manifold-based models.

Next is a synopsis of the paper. In Sect. 2 we recall basic properties of oblivious ε-subspace embeddings and their use in the so-called randomized QR factorization. Section 3 reports the derivation of the $S^\top S$-SVD, its properties along with a fast algorithm for its computation. Some important scenarios where the $S^\top S$-SVD can be adopted are explored in Sect. 3.1 whereas in Sect. 4 we employ the $S^\top S$-SVD to bound the distance of $S^\top S$-orthogonal matrices from “standard” orthogonality. The sketched version of (1) is presented and analyzed in Sect. 5. The paper ends with some conclusions in Sect. 6.

All the experiments reported in this paper have been run using Matlab (version 2024b) on a machine with a 1.2GHz Intel quad-core i7 processor with 16GB RAM on an Ubuntu 20.04.2 LTS operating system.

Notation. We use $x^\top$ to denote the transpose of the vector x. We use the Euclidean norm for vectors, and the associated induced matrix norm for matrices, which we call 2-norm $\|\cdot\|_2$. We also report results using the Frobenius matrix norm, denoted by $\|\cdot\|_F$. Given a matrix A, we denote with $A^\dagger$ its pseudo-inverse (Moore-Penrose inverse). Exact arithmetic is assumed throughout.

With some abuse of notation, throughout the paper we shall adopt the terminology “orthogonal matrix” also for tall rectangular matrices whose columns are orthonormal, without requiring the matrix to be square.

2 Subspace embeddings and randomized QR

In the last decade, randomization-based tools have been shown to be an important aid in decreasing the computational cost of a number of algorithms in numerical linear algebra and scientific computing in general.

One such tool is given by (oblivious) ε-subspace embeddings; see, e.g., [25, Sect. 8.7]. Here we use the notation adopted in [3].

Definition 1 ([3, Definition 2.3]) Let $\mathcal{V}$ be a k-dimensional subspace of $\mathbb{R}^m$. Then, given $\varepsilon \in (0, 1)$ and $s \le m$, a linear map $S \in \mathbb{R}^{s\times m}$ is said to be an (ε, δ, k)-subspace embedding for $\mathcal{V}$ if

$$(1-\varepsilon)\|v\|^2 \le \|Sv\|^2 \le (1+\varepsilon)\|v\|^2, \quad \text{for any } v \in \mathcal{V}, \tag{2}$$

holds with probability $1-\delta$, at least.

We observe that from a deterministic point of view, $S^\top S$ defines a *semi*-norm, whereas within a probabilistic setting, the definition above ensures that $\|Sv\|^2$ is greater than zero for $v \in \mathcal{V}$, $v \ne 0$, with high probability; see, e.g., [4, Proposition 3.3]. A fact following from (2) is that with high probability it holds that

$$-\varepsilon\|v\|^2 \le v^\top(I - S^\top S)v \le \varepsilon\|v\|^2, \tag{3}$$

which measures the distance of $S^\top S$ from acting as the identity.

Various choices of randomized linear maps S have been proposed in the literature. An incomplete list includes Gaussian transformations, sparse sign matrices, and subsampled trigonometric functions; see, e.g., [25, Sect. 9]. Given the dimension k of the subspace to be embedded, the threshold ε, and the failure probability δ, one can select suitable sketching dimensions s to ensure (2). Our derivations do not depend on the nature of the sketching S as long as (2) holds. Hence, in the following we will only assume that the adopted S is a randomized $(\varepsilon, \delta, \mathrm{rank}(A))$-subspace embedding for $\mathrm{Range}(A)$ so that the $S^\top S$-norm is well-defined on this space with high probability.

Definition 2 Given a matrix $P \in \mathbb{R}^{m\times n}$ and a sketching matrix $S \in \mathbb{R}^{s\times m}$, if S is a randomized $(\varepsilon, \delta, \mathrm{rank}(P))$-subspace embedding for $\mathrm{Range}(P)$ then we can define the following matrix norms

$$\|P\|_{S,F}^2 := \mathrm{trace}(P^\top S^\top S P),$$

and

$$\|P\|_{S,2}^2 := \max_{\|x\|_2=1} x^\top P^\top S^\top S P x.$$

Thanks to (2), it follows that

$$(1-\varepsilon)\|P\|_F^2 \le \|P\|_{S,F}^2 \le (1+\varepsilon)\|P\|_F^2, \quad \text{and} \quad (1-\varepsilon)\|P\|_2^2 \le \|P\|_{S,2}^2 \le (1+\varepsilon)\|P\|_2^2;$$

see, e.g., [20, Corollary 3.1]. Note that both results also readily follow from using the singular value bounds in Theorem 3 below, derived in [13].

Thanks to their ability in preserving norms up to a small distortion parameter, oblivious ε-subspace embeddings are the backbone of randomized algorithms for QR factorizations; see, e.g., [3, 15, 20, 23]. These procedures are able to cut down the cost of computing QR factorizations. On the other hand, they provide factorizations of the form $A = QR$ with *almost* orthogonal Q factor; we refer to sect. 4 for a more detailed discussion of this distance to orthogonality.

In the following we shall say that Q is $S^\top S$-orthogonal, with high probability, if $Q^\top S^\top S Q = I$, while Q is orthogonal if $Q^\top Q = I$, where I is the identity matrix of dimension equal to the number of columns of Q.

One of the main advantages in employing the sketching matrix S is the decrease in the cost of the inner products, as vectors of length s – instead of m – are involved. On the other hand, S needs to be applied to each of the n columns of A. Therefore, matrix-vector products with S must be as cheap as possible to obtain an effective randomized orthogonalization; see, e.g., the detailed analysis in [3, Sect. 2.4]. For instance, when S amounts to a subsampled trigonometric transformation, like in Example 2– 3 below, performing Sx costs $\mathcal{O}(m \log s)$ floating point operations (flops) if the action of S is judiciously implemented [32]. This means that only $\mathcal{O}(mn \log s)$ flops are needed to compute SA. Notice that these costs can be further reduced in case of a sparse A.

Most available analyses on sketched QR procedures focus on backward stability issues, by estimating the residual norm $\|A - QR\|_*$, and controlling the growth of the condition number of Q, $\kappa_*(Q) = \|Q\|_* \|Q^\dagger\|_*$, for $* = 2, F$; see, e.g., [3, 15]. We enrich this analysis by investigating how far the columns of an $S^\top S$-orthogonal matrix are from an orthogonal matrix. If $S^\top S$ were positive definite this investigation would be very short, as this distance would depend on the condition number of $S^\top S$. Our probabilistic setting requires extra work. To give a glimpse of the type of expected results, we report an upper bound on the cosine of the angle between two different columns of an $S^\top S$-orthogonal matrix P, and thus the Q factor computed by the sketched QR decomposition. To this end, we first recall a result from [9], where we use the notation

$$\cos \angle(u, v) = \frac{u^\top v}{\|u\|_2 \, \|v\|_2},$$

for the cosine of the angle between two nonzero vectors.

Lemma 1 ([9, Lemma 4.3]) *Let $S \in \mathbb{R}^{s \times m}$ be such that (2) holds for two vectors $u, v \in \mathbb{R}^m$ and $\varepsilon > 0$. Then with high probability it holds that*

$$\frac{\cos \angle(u, v) - \varepsilon}{1 + \varepsilon} \leq \cos \angle(Su, Sv) \leq \frac{\cos \angle(u, v) + \varepsilon}{1 - \varepsilon}.$$

Now, let $P = [p_1, \ldots, p_n]$ be such that $(Sp_i)^\top (Sp_j) = \delta_{ij}$, where δ_{ij} denotes the Kronecker delta. Then Lemma 1 implies that

$$i \neq j, \;\; (Sp_i)^\top (Sp_j) = 0 \;\Rightarrow\; |\cos \angle(p_i, p_j)| \leq \varepsilon.$$

Therefore, the value ε provides an upper bound also on the distortion of the cosine of the angle between two columns of an $S^\top S$-orthogonal matrix. More precise bounds on the distance from Euclidean orthogonality are given in sect. 4. To this end, in the following section we discuss the $S^\top S$-SVD factorization, which will be used as a strategic tool for measuring the distance between subspaces, in agreement with the role of the standard SVD [14, Sect. 2.5].

3 The $S^\top S$-SVD

In [30, Theorem 3], Van Loan derived a generalization of the SVD of a given matrix in terms of nonstandard inner products.

Definition 3 ([30, Definition 3]) Given $A \in \mathbb{R}^{m\times n}$, $m \geq n$, and a symmetric and positive definite matrix $\mathcal{H} \in \mathbb{R}^{m\times m}$, the $\mathcal{H}$, 2-singular values of A are the elements of the following set

$$\mu_{\mathcal{H},2}(A) := \left\{\mu \geq 0 \text{ s. t. } \mu \text{ is a stationary value of } \frac{\|Ax\|_{\mathcal{H}}}{\|x\|_2},\ x \neq 0\right\}. \tag{4}$$

In our context, $S^\top S$ is not positive definite from a deterministic viewpoint. Nonetheless, we can still formalize an SVD where the left singular vectors are $S^\top S$-orthogonal, and the factorization holds with high probability.

Theorem 1 *Let $A \in \mathbb{R}^{m\times n}$, $m \geq n$, and let $S \in \mathbb{R}^{s\times m}$ be a randomized $(\varepsilon, \delta, rank(A))$-subspace embedding of $Range(A)$, $s \geq rank(A)$. Then there exist an $S^\top S$-orthogonal matrix $W \in \mathbb{R}^{m\times r}$, an orthogonal matrix $V \in \mathbb{R}^{n\times r}$ with $r = \min\{s, n\}$, and a diagonal matrix $\Theta = \operatorname{diag}(\theta_1, \ldots, \theta_r)$ with $\theta_1 \geq \ldots \geq \theta_r \geq 0$ such that A can be written as*

$$A = W\Theta V^\top, \tag{5}$$

where the equality holds with probability $1 - \delta$.

Proof The result directly comes from [30, Theorem 3] as long as (2) holds true. Indeed, in this case $S^\top S$ is positive definite on $\text{Range}(A)$. On the other hand, the property (2) holds with probability $1 - \delta$ which means that the same happens for our factorization (5). □

The next corollary reports a first characterization of the $S^T S$-singular values θ_i's; its proof follows the steps of that of Theorem 3 in [30].

Corollary 1 *The nonegative scalars θ_i, $i = 1, \ldots, r$, in Theorem 1 are stationary values of the function $\mu(x) = \|Ax\|_{S^\top S}/\|x\|_2$ with $x \neq 0$.*

Proof We consider the equivalent functional $\mu(x) = \|Ax\|^2_{S^\top S}$, with the constraint $\|x\|_2 = 1$, and define the Lagrangian $L(x, \lambda) = \|Ax\|^2_{S^\top S} - \lambda(\|x\|_2^2 - 1) = x^\top A^\top S^\top S A x - \lambda(x^\top x - 1)$. Stationary points are determined by taking derivatives, and they are the zeros of $\det(A^\top S^\top S A - \mu^2 I) = 0$.

Let $SA = U_1\widehat{\Theta}V^\top$, $U_1 \in \mathbb{R}^{s\times r}$, $\widehat{\Theta} \in \mathbb{R}^{r\times n}$, $V \in \mathbb{R}^{n\times n}$, denote the standard SVD of SA. If $s \geq n$, then $r = n$ and $\widehat{\Theta} = \Theta$. Otherwise, $r = s$ and $\widehat{\Theta} = [\Theta, 0]$. Therefore, $A^\top S^\top SA = V\widehat{\Theta}^2 V^T$, where $\widehat{\Theta}^2 = \mathrm{diag}(\Theta^2, 0)$. Hence

$$\det(A^\top S^\top SA - \mu^2 I) = \det(\widehat{\Theta}^2 - \mu^2 I) = \prod_{i=1}^{n}(\widehat{\Theta}^2_{i,i} - \mu^2).$$

so that the first r roots coincide with the θ_i's. □

Returning to Theorem 1, we notice that since S is a subspace embedding for Range(A), s will always be required to be larger than the rank of A, even if this is less than n. The next remark highlights the implications of possible null θ_is.

Remark 1 Since $S \in \mathbb{R}^{s\times m}$ is asked to be a randomized $(\varepsilon, \delta, \mathrm{rank}(A))$-subspace embedding of Range(A), the sketching dimension s should be selected in terms of rank(A). For simplicity, n, the number of columns of A, can be used instead of rank(A). While this strategy certainly works, it is important to stress that it can lead to unnecessarily large values of s, especially in case of a low-rank A, which can significantly hinder the computational performance of the approach. We refer to Example 1 below for an example with a low-rank square matrix.

The factors in the $S^\top S$-SVD of Theorem 1 can be computed by means of the following steps:

1. Compute the SVD of $SA = U_1\Theta V^\top$
2. Define $W := A(\Theta V^\top)^\dagger$

Next, a simple Matlab implementation for A full rank and equivalent to the procedure above is reported.

```
[~,R]=qr(S*A,0);
[~,Theta,V]=svd(R);
W=A*(V/Theta);
```

Notice that in the commands above the only operations whose cost depends on the dimensionality m of the problem are the application of the sketching SA and the retrieval of the left singular vectors $W = AV\Theta^\dagger$. Both these steps can take advantage of the possible sparsity of A. In all reported experiments, we computed the $S^\top S$-SVD by the above commands.

This procedure is closely related to the approach derived in [13], where, however, the authors restricted the use of the term "sketched SVD" of A to the SVD of $SA \in \mathbb{R}^{s\times n}$. With our approach, instead, the left singular matrix W yields an $S^\top S$-orthogonal basis for the range of A. We thus refrain from using the adjective "sketched" for (5), as the latter refers to a decomposition of the whole matrix A.

Remark 2 If W is $S^\top S$-orthogonal and U is square in the Euclidean sense, then WU is still $S^\top S$-orthogonal. The same holds if U is tall and orthogonal, although $S^\top S$-orthogonality will be in terms of the number of columns of U.

Remark 3 The $S^\top S$-SVD of A can also be determined by means of the randomized QR factorization of Sect. 2: compute $A = QR$, with $Q^\top S^\top SQ = I$. Then, compute the standard SVD of R, namely $R = U\Theta V^\top$, so that the $S^\top S$-SVD of A is given by

$$A = W\Theta V^\top, \quad W := QU.$$

Using the randomized QR, this implementation of the $S^\top S$-SVD may turn out to be more robust than the one illustrated above which could be prone to numerical instabilities due to the computation of $\Theta^\dagger$. On the other hand, the former is in general more expensive due to the computation of Q. We also stress the *single-pass* nature of this implementation of the $S^\top S$-SVD, which allows us to access A only once, as required by certain data streaming models. At the same time, the $S^\top S$-SVD is still able to compute a decomposition with ($S^\top S$-)orthogonal factors, unlike other state-of-the-art randomized methods; see, e.g., the discussion in [18, Sect. 5.5].

By employing the $S^\top S$-SVD we can provide alternative definitions of the $S^\top S$-matrix norms in Definition 2, that is

$$\|A\|_{S,F}^2 = \sum_{i=1}^n \theta_i^2, \quad \text{and} \quad \|A\|_{S,2} = \theta_1.$$

As its standard counterpart, also the $S^\top S$-SVD and the $S^\top S$-singular values θ_i's fulfill certain optimality conditions as long as the latter are formulated in the right norm i.e., the $S^\top S$-norm.

Theorem 2 *Let $A = W\Theta V^\top$ with $\Theta = diag(\theta_1, \ldots, \theta_n)$ be the $S^\top S$-SVD of $A \in \mathbb{R}^{m\times n}$, $m \geq n$. Let $W_k \in \mathbb{R}^{m\times k}$, $V_k \in \mathbb{R}^{n\times k}$ collect the first k columns of W and V, respectively, and Θ_k be the square top left leading part of Θ. Let $S \in \mathbb{R}^{s\times m}$ be a randomized (ε, δ, n)-subspace embedding of any subspace of $\mathbb{R}^m$ of dimension n. Then with high probability*

$$\min_{\substack{B\in\mathbb{R}^{m\times n}\\ rank(B)=k}} \|A-B\|_{S,F}^2 = \sum_{i=k+1}^n \theta_i^2, \qquad W_k\Theta_k V_k^\top = \arg\min_{\substack{B\in\mathbb{R}^{m\times n}\\ rank(B)=k}} \|A-B\|_{S,F}, \tag{6}$$

$$\min_{\substack{B\in\mathbb{R}^{m\times n}\\ rank(B)=k}} \|A-B\|_{S,2}^2 = \theta_{k+1}, \qquad W_k\Theta_k V_k^\top = \arg\min_{\substack{B\in\mathbb{R}^{m\times n}\\ rank(B)=k}} \|A-B\|_{S,2}, \tag{7}$$

and

$$\theta_k = \max_{\mathcal{U}, dim(\mathcal{U})=k} \min_{\substack{x\in\mathcal{U}\\ \|x\|_2=1}} \|Ax\|_{S^\top S} = \min_{\mathcal{U}, dim(\mathcal{U})=n-k+1} \max_{\substack{x\in\mathcal{U}\\ \|x\|_2=1}} \|Ax\|_{S^\top S}. \tag{8}$$

Proof All these optimality results can be shown by mimicking the proofs of the corresponding results for standard SVD and recalling that the left $S^\top S$-singular vectors are $S^\top S$-orthogonal. □

In the analysis that follows, it will be important to relate the $S^\top S$-singular values of A to their standard counterparts. Since S is a randomized $(\varepsilon, \delta, \mathrm{rank}(A))$-subspace embedding for Range(A), we know that

$$(1-\varepsilon)\|Ax\|_2^2 \le \|SAx\|_2^2 \le (1+\varepsilon)\|Ax\|_2^2,$$

so that

$$(1-\varepsilon)\sigma_n^2 \le \|SAx\|_2^2 \le (1+\varepsilon)\sigma_1^2,$$

where σ_1 and σ_n are the largest and smallest (standard) singular values of A, respectively. This means that all the $S^\top S$-singular values of A are included in the interval $[\sqrt{1-\varepsilon}\cdot\sigma_n, \sqrt{1+\varepsilon}\cdot\sigma_1]$. More accurate bounds can be obtained by monitoring the change in each singular value. Such an analysis, together with the corresponding bounds for the right singular vectors was carried out in [13], and since then it has been rediscovered a few times, also with different proofs, see, e.g., [16, Theorem 2.2].

Theorem 3 ([13]) *Let S be a randomized $(\varepsilon, \delta, rank(A))$-subspace embedding for Range(A), and let $A = W\Theta V^\top$ be the $S^\top S$-SVD of $A \in \mathbb{R}^{m\times n}$, $m \ge n$, and $A = U\Sigma Y^\top$, $\Sigma = diag(\sigma_1, \ldots, \sigma_n)$, its standard SVD. Then with high probability,*

$$\sqrt{1-\varepsilon}\cdot\sigma_k \le \theta_k \le \sqrt{1+\varepsilon}\cdot\sigma_k, \quad \textit{for all } k = 1, \ldots, n. \tag{9}$$

Theorem 3 is an important asset in the analysis of sketched numerical procedures. For instance, the rank of A can be read by looking at the $S^\top S$-singular values. Indeed, $\theta_k = 0$ if and only if $\sigma_k = 0$, with high probability. Moreover, the *truncated* $S^\top S$-SVD $W_k\Theta_k V_k^\top$ in (6) satisfies

$$\left\|A - W_k\Theta_k V_k^\top\right\|_{S,F}^2 = \sum_{i=k+1}^{n} \theta_i^2 \le (1+\epsilon)\sum_{i=k+1}^{n}\sigma_i^2, \text{ and}$$
$$\left\|A - W_k\Theta_k V_k^\top\right\|_{S,2} = \theta_{k+1} \le \sqrt{1+\epsilon}\cdot\sigma_{k+1}.$$

Therefore, the truncated $S^\top S$-SVD, which is optimal in the $S^\top S$-norm, attains an error that with high probability is a small multiple of the optimal error in the Frobenius and Euclidean norm. This opens to the design of fast, suboptimal low-rank approximations with $S^\top S$-orthogonal factors.

It is also interesting to compare the $S^\top S$-SVD with the randomized SVD [18], which relies on a randomized range finder algorithm in place of subspace embeddings. Given a target rank k and an oversampling parameter ℓ, the range finder first draws a sketching matrix $\Omega \in \mathbb{R}^{n\times(k+\ell)}$, which does not necessarily fulfill any subspace embedding properties, and it then constructs a matrix $Q \in \mathbb{R}^{m\times(k+\ell)}$ whose columns form an orthonormal basis of Range$(A\Omega)$. The SVD-like approximation provided by the randomized SVD algorithm is $\widehat{A} := (Q\widehat{U})\widehat{\Sigma}\widehat{V}^\top \approx A$ where $\widehat{U}\widehat{\Sigma}\widehat{V}^\top$ is the standard SVD of $Q^\top A$. By partitioning the right singular vectors of A as $V = [V_1, V_2]$, $V_1 \in \mathbb{R}^{n\times k}$, $V_2 \in \mathbb{R}^{n\times(n-k)}$ and by denoting $\Omega_1 := V_1^\top\Omega$, $\Omega_2 := V_2^\top\Omega$, under the assumptions $\mathrm{rank}(\Omega_1) = k$ and $\gamma_k := \sigma_{k+1}/\sigma_k < 1$, in [27, Theorem 9] it was shown that[1]

$$\frac{\sigma_j}{\sqrt{1+\gamma_j^2\|\Omega_2\Omega_1^\dagger\|_2^2}} \le \widehat{\sigma}_j \le \sigma_j, \quad \text{for all } j = 1, \ldots, k,$$

where $\widehat{\sigma}_j$ denotes the jth singular values of $\widehat{A}$. The width of the interval containing the ratio $\widehat{\sigma}_j/\sigma_j$ depends on the gap between the corresponding consecutive singular values of A and on the quality of the random matrix Ω. In our case, the corresponding ratio for the $S^\top S$-singular value θ_j lies in the interval $[\sqrt{1-\varepsilon}, \sqrt{1+\varepsilon}]$, which is related to the considered probabilistic confidence. A more detailed comparison is postponed to later research.

3.1 On the applicability of the $S^\top S$-SVD

In many applications the singular values of the given matrix A are characterized by particular decay properties or by the presence of interval gaps. In the latter case, for instance, the gap may correspond to a very ill-conditioned A; truncation procedures have been classically adopted to eliminate what can be recognized as noisy data, see, e.g., [31]. In these scenarios, the $S^\top S$-singular values θ_is of A can be computed at low cost, and analyzed in place of the original ones. Thanks to (9), the qualitative behavior of the θ_is will be the same as that of the singular values σ_is of A. For example, a distribution gap of the σ_is can be captured by inspecting a similar gap in the θ_is (see Figure 1 below), and thus determining the number of singular values to be retained. Such a strategy can be employed to determine the target rank in probabilistic or deterministic procedures that seek a low rank approximation of A of given rank, such as truncated SVD, CUR, and Nyström type approximations [8, 18, 25]. Moreover, the $S^\top S$-singular values θ_is can be exploited to determine a priori the truncation threshold of the solution in regularized linear least squares problems associated with ill-posed inverse problems; see, e.g., [19].

The next example illustrates that the $S^\top S$-singular values are indeed able to capture the low numerical rank occurring in the original data.

Example 1 We consider the Cauchy matrix $C \in \mathbb{R}^{n\times n}$, $n = 5\,000$, defined entry-wise as $C_{i,j} = 1/(x_i + y_j)$ where x_i and y_j are taken as the ith and jth nodes of a grid of n equidistant points in [2, 100] and $[-1\,000, -500]$, respectively. Cauchy matrices present a fast decay in their singular values if, e.g., the values $\{x_i\}$ and $\{y_j\}$ come from disjoint sets; see, e.g., [5, Sect. 4.1]. Thus, the considered matrix C is numerically low-rank. We show that the $S^\top S$-SVD amounts to a practical, efficient way for estimating its numerical rank.

As sketching $S \in \mathbb{R}^{s\times m}$ we consider the following subsampled trigonometric transformation

$$S = \sqrt{\frac{m}{s}} DFE, \tag{10}$$

[1] Note that in the original version of [27, Theorem 9] the term γ_j^{4q+2} appears in place of γ_j^2 as we report. The scalar q consists of the number of performed subspace iterations. Here we assume $q = 0$, namely the plain randomized SVD algorithm is adopted.

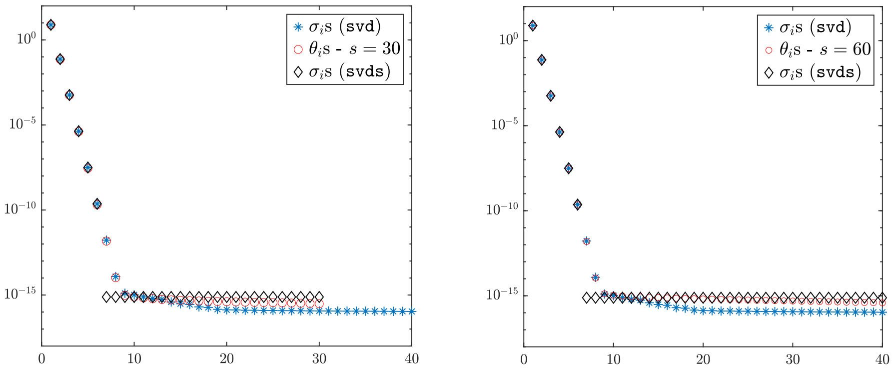

Fig. 1 Example 1. First 40 standard singular values σ_is computed by svd (blue stars), first ℓ $S^\top S$-singular values θ_i's (red circles), and first ℓ standard singular values σ_is computed by svds (black diamond) for different values of the sketching dimension s and $\ell = \min\{s, 40\}$. Left: $s = 30$. Right: $s = 60$. The reported results have been averaged over 50 runs

where $E \in \mathbb{R}^{m\times m}$ is a diagonal matrix with Rademacher entries (i.e., the diagonal entries are randomly chosen as ± 1 with equal probability), $D \in \mathbb{R}^{s\times m}$ contains s randomly selected rows of the identity matrix, and $F \in \mathbb{R}^{m\times m}$ is the discrete cosine transform. Due to our rather basic implementation, the cost of applying the matrix S to an $m \times n$ matrix is $\mathcal{O}(mn \log m)$ operations[2].

It was shown in [29] that if $s = \mathcal{O}(\varepsilon^{-2}(k + \log \frac{m}{\delta}) \log \frac{k}{\delta})$, then (10) is an oblivious ε-subspace embedding for any k-dimensional subspace of $\mathbb{R}^m$. Nonetheless, numerical evidence suggests that selecting the smaller sketching dimension $s = \mathcal{O}(\varepsilon^{-2}\frac{k}{\delta})$ works well in practice; see, e.g., [25, Sect. 9].

For the given C, although it is known that the rank is low, the actual dimension of Range(C) is not known a-priori. We thus select rather moderate values of s and in Figure 1 we report the first 40 singular values σ_is of C computed by svd (blue stars), the first $\ell := \min\{s, 40\}$ $S^\top S$-singular values θ_is (red circles), and the first ℓ standard singular values computed by svds (black diamonds) for $s = 30$ (left) and $s = 60$ (right).

Figure 1 shows that the $S^\top S$-singular values are able to match rather well the true singular values of A, especially those larger than 10^{-15}, for both the tested values of s. The $S^\top S$-SVD thus perfectly encodes the (standard) numerical rank of C. This does not really happen when running the memory-saving deterministic Matlab function svds[3]. Indeed, this routine is unable to catch $\sigma_7 \approx 10^{-12}$ and $\sigma_8 \approx 10^{-14}$.

The computation of the $S^\top S$-singular values required only 0.14 and 0.15 seconds for $s = 30$ and $s = 60$, respectively. These timings are one order of magnitude smaller than the running time devoted to the computation of all the standard singular values of[4] C, which amounts to 3.31 seconds. The matlab function svds required 1.82 ($s = 30$) and 3.6 ($s = 60$) seconds to compute the first s σ_i's. In addition to being still larger

[2] The action of F is computed via the matlab function dct.

[3] This is still the case when changing the default setting of svds to, e.g., `sigma=svds(A,s,'largest','Tolerance', 1e-15,'MaxIterations',2e3)`.

[4] `sigma=svd(C)`.

than the running time required by our approach, these numbers show that the cost of `svds` grows linearly with the parameter s. On the other hand, the cost of our approach is practically insensitive to s.

We conclude by mentioning that also the randomized SVD algorithm with a subsampled trigonometric sketching matrix $\Omega \in \mathbb{R}^{n\times(k+\ell)}$, $k = s$, $\ell = 5$, is able to match all the first eight singular values of A with running times that are slightly larger than the ones attained by the $S^\top S$-SVD. This is mainly due to the explicit projection $Q^\top A$ performed by the former scheme. ◇

4 Orthogonality properties

In using the sketched approach, a natural question is how far an $S^\top S$-orthogonal matrix is from being orthogonal in the Frobenius and Euclidean sense. The $S^\top S$-singular values can be employed for these purposes. As a counterpart result, we will also derive bounds on the distance of standard orthogonal matrices from $S^\top S$-orthogonality.

Proposition 1 *Let $P \in \mathbb{R}^{m\times n}$ have $S^\top S$-orthonormal columns with S being a randomized $(\varepsilon, \delta, rank(P))$-subspace embedding for Range(P). Then, with high probability,*

$$\|P^\top P - I\|_F \le \frac{\varepsilon}{1-\varepsilon}\sqrt{n}, \quad and \quad \|P^\top P - I\|_2 \le \frac{\varepsilon}{1-\varepsilon}.$$

Proof Since P has $S^\top S$-orthonormal columns, all $S^\top S$-singular values θ_j of P are equal to 1. Now, let $P = U\Sigma V^\top$, $\Sigma = \text{diag}(\sigma_1, \ldots, \sigma_n)$, be the standard SVD of P. Since $\sigma_j^2 \le 1/(1-\varepsilon)$ for all $j = 1, \ldots, n$, it holds

$$\|P^\top P - I\|_2 = \|\Sigma^2 - I\|_2 = \max_{i=1,\ldots,n} |\sigma_i^2 - 1| \le \frac{1}{1-\varepsilon} - 1 = \frac{\varepsilon}{1-\varepsilon}.$$

Similarly,

$$\|P^\top P - I\|_F^2 = \|\Sigma^2 - I\|_F^2 = \sum_{i=1}^n (\sigma_i^2 - 1)^2 \le \sum_{i=1}^n \left(\frac{1}{1-\varepsilon} - 1\right)^2 = n\left(\frac{\varepsilon}{1-\varepsilon}\right)^2.$$

□

The bounds of the above proposition can be proved in different ways. For example, directly using the fundamental property in (3), the orthogonality $P^\top S^\top SP = I$ and the bounds on the 2-norm of P, we can obtain the bounds of $\|P^\top P - I\|_2$ by estimating the inner product $x^\top(P^\top P - I)x = x^\top P^\top(I - S^\top S)Px$.

We also observe that the previous bounds hold for any $S^\top S$-orthogonal matrix, thus also for the one stemming from the randomized (sketched) QR decomposition [3].

Proposition 2 *Let $T \in \mathbb{R}^{m\times n}$ have orthonormal columns. Let S be a randomized $(\varepsilon, \delta, rank(T))$-subspace embedding for Range(T). Then, with high probability,*

$$\|T^\top S^\top ST - I\|_F \le \varepsilon\sqrt{n}, \quad and \quad \|T^\top S^\top ST - I\|_2 \le \varepsilon.$$

Table 1 Example 2. Loss of orthogonality (Frobenius and Euclidean norm) of the first n left $S^\top S$-singular left vectors of A (i.e. the columns of W) as s varies. "Time" (in secs) is the running time to compute the $S^\top S$-SVD of A. The reported results are averaged over 50 runs

s	$\|W^\top W - I\|_F$	$\|W^\top W - I\|_2$	Time (s)
$55 \log n$	16.84	0.99	0.27
$60 \log n$	17.04	0.99	0.28
$65 \log n$	17.13	0.99	0.28

Proof Let $T = W\Theta V^\top$ be the $S^\top S$-singular value decomposition of T, and recall that the standard singular values σ_i of T are all equal to 1. Then from $\|T^T S^\top ST - I\|_* = \|\Theta^2 - I\|_*$ and using Theorem 3, it follows

$$\|T^\top S^\top ST - I\|_2 \leq \max|\theta_i^2 - 1| \leq \max|(1+\varepsilon)\sigma_i^2 - 1| = \varepsilon,$$

while

$$\|T^\top S^\top ST - I\|_F^2 \leq \sum_i |\theta_i^2 - 1|^2 \leq \sum_i |(1+\varepsilon)\sigma_i^2 - 1|^2 = \varepsilon^2 n,$$

from which both results follow. □

Again, using the property in (3), the orthogonality bound $\|T^\top S^\top ST - I\|_2 \leq \varepsilon$ could also be proved by noticing that $\|T^\top S^\top ST - I\|_2 = \|T^\top S^\top ST - T^\top T\|_2$ and that $x^\top T^\top(S^\top S - I)Tx = y^\top(S^\top S - I)y$ with $y = Tx$, where $\|y\| = \|x\|$.

In the next example we illustrate what to expect in practice, in terms of distance from full orthogonality, when relying on results such as that in Proposition 1.

Example 2 We consider the matrix $A \in \mathbb{R}^{m\times n}$, $m = 300\,000$, $n = 300$, generated by the matlab function `A=sprand(m,n,0.003,1e-10)`. Therefore, only about 0.3% of the entries of A are non-null, and the nonzero values are taken from a uniform random distribution. Moreover, $\kappa(A) = 10^{10}$.

For this example we change the nature of the sketching and consider $S \in \mathbb{R}^{s\times m}$ to be a Gaussian matrix. The cost of computing SA with $A \in \mathbb{R}^{m\times n}$ is now $\mathcal{O}(s \cdot \mathrm{nnz}(A))$ flops where $\mathrm{nnz}(A)$ denotes the number of nonzero entries of A. Moreover, to embed an n-dimensional subspace with failure probability δ, it is sufficient to choose $s = \mathcal{O}(\varepsilon^{-2} \log n \log \frac{1}{\delta})$; see, e.g., [28, Theorem 2].

We are interested in illustrating the bounds in Proposition 1. In Table 1, for different values of s we report $\|W^\top W - I\|_*$, $* = 2, F$, where the columns of $W \in \mathbb{R}^{m\times n}$ are the first n left $S^\top S$-singular vectors of A. If we set $\varepsilon = 0.5$, as it is common when working with oblivious subspace embeddings, then $\varepsilon/(1-\varepsilon) = 1$ so that the bounds in Proposition 1 are satisfied whenever $\|W^\top W - I\|_F \leq \sqrt{300} \leq 18$ and $\|W^\top W - I\|_2 \leq 1$.

For $\varepsilon = 0.5$ and $\delta = 10^{-6}$, we should select $s = \mathcal{O}(56 \log n)$ for S to be a randomized (ε, δ, n)-subspace embedding with probability $1 - \delta$. From the results in Table 1 we can see that, with this set of parameters, the bounds in Proposition 1 are always fulfilled. In Table 1 we also document the running times devoted to computing the $S^\top S$-SVD of A. We notice that these timings only mildly depend on the sketching dimension s, and they are rather competitive with the timing required to compute the standard SVD of A, which is 2.89 seconds with the matlab `svd` function[5]. This is mainly due to the need of `svd` to work with a matrix allocated in full format[6]. For the chosen sketching this drawback does not affect the $S^\top S$-SVD routine, that can fully exploit the sparsity of A. ◇

5 The nearest $S^\top S$-orthogonal matrix

With the goal of reducing the solution cost when dealing with large dimensional matrices, in [7] the authors derived an iterative algorithm for (1). With the same purposes, in this section we propose to explore the use of sketching techniques. To this end, we first transform the problem by using the $S^\top S$-norm. Given a matrix $A \in \mathbb{R}^{m\times n}$, $m \geq n$, we consider the following matrix minimization problem

$$\min_{\substack{Q\in\mathbb{R}^{m\times n}\\ Q^\top S^\top SQ=I}} \|A - Q\|_{S,*}, \quad * = 2, F. \tag{11}$$

Thanks to the $S^\top S$-SVD in Sect. 3, we are going to show the existence and uniqueness of its solution and relate the latter to the solution of the original problem (1).

Now let $A = W\Theta V^\top$ be the $S^\top S$-SVD of A. Then

$$P = WV^\top, \tag{12}$$

is still $S^\top S$-orthogonal thanks to Remark 2. Moreover, P is the orthogonal factor of the *randomized* polar decomposition of A. Indeed, we can write

$$A = PH, \quad \text{with } H = V\Theta V^\top, \tag{13}$$

with $\Theta = \text{diag}(\theta_1, \ldots, \theta_n)$ and $\theta_i \geq 0$, so that H is positive semidefinite. Notice that, to the best of our knowledge, a randomization-based polar decomposition has never been proposed prior this work. Indeed, the decomposition (13) is a natural consequence of our $S^\top S$-SVD thanks to the availability of the left singular matrix W.

In the next theorem we prove that P defined in (12) solves the problem (11). To this end, we notice that we can parametrize all $S^\top S$-orthogonal matrices spanning $\text{Range}(A)$ by introducing the following set

$$\mathcal{Q}_S(A) = \{Q \in \mathbb{R}^{m\times n} \,:\, Q = WLV^\top, \text{ with } L^\top L = I,\ A = W\Theta V^\top\}.$$

The following result holds.

[5] `[U, Sigma, V] = svd(A, 0)`.

[6] The time required for changing A to full format was not included.

Theorem 4 *Let $A = W\Theta V^\top$ be the $S^\top S$-SVD of A where S is a randomized $(\varepsilon, \delta, rank(A))$-subspace embedding of $Range(A)$. Let $A = PH$ with $P \in \mathcal{Q}_S(A)$ defined in (12) and $H = V\Theta V^\top$. Then with high probability,*

$$\|A - P\|_{S,*} = \min_{Q \in \mathcal{Q}_S(A)} \|A - Q\|_{S,*}, \tag{14}$$

where $* = 2, F$.

Proof Using $A = PH$ and the $S^\top S$-orthogonality of P, we first notice that

$$\|A - P\|_{S,*}^2 = \|P(H - I)\|_{S,*}^2 = \|H - I\|_*^2 = \|\Theta - I\|_*^2.$$

We first focus on $* = F$. For the same reasonings as above, for any $Q \in \mathcal{Q}_S(A)$ and using that $L^\top L = I$, we have

$$\begin{aligned}
\|A - Q\|_{S,F}^2 &= \|\Theta - L\|_F^2 = \|\Theta - I + I - L\|_F^2 \\
&= \text{trace}\left((\Theta - I)^2 + (\Theta - I)(I - L) + (I - L)^\top(\Theta - I) + (I - L)^\top(I - L)\right) \\
&= \text{trace}\left((\Theta - I)^2 + \Theta(I - L) + (I - L)^\top\Theta\right) \\
&= \|\Theta - I\|_F^2 + 2\sum_{i=1}^n \theta_i(1 - \ell_{ii}).
\end{aligned}$$

The orthogonality of L also implies that $\ell_{i,i} \le 1$ with the equality attained for all is if and only if $L = I$, i.e., $Q = P$. Therefore, $2\sum_{i=1}^n \theta_i(1 - \ell_{ii}) \ge 0$ for $L \ne I$. In conclusion, we thus have

$$\|A - Q\|_{S,F}^2 \ge \|\Theta - I\|_F^2 = \|A - P\|_{S,F}^2 ,$$

which shows the first result.

For $* = 2$, we have that

$$\|A - P\|_{S,2}^2 = \max_{\|x\|_2 = 1} x^\top(\Theta - I)^2 x = \max_{i=1,\dots,n} (\theta_i - 1)^2,$$

which means that the vector x attaining the maximum is one of the vectors of the canonical basis of $\mathbb{R}^n$, namely an $e_{\bar{p}}$ for a certain $\bar{p} \in \{1, \dots, n\}$.

For a generic $S^\top S$-orthogonal matrix $Q = WLV^T$ it holds

$$\begin{aligned}
\|A - Q\|_{S,2}^2 &= \|\Theta - L\|_2^2 = \|\Theta - I + I - L\|_2^2 \\
&= \max_{\|x\|_2 = 1} x^\top\left((\Theta - I)^2 + \Theta(I - L) + (I - L)^\top\Theta\right) x \\
&\ge e_{\bar{p}}^\top\left((\Theta - I)^2 + \Theta(I - L) + (I - L)^\top\Theta\right) e_{\bar{p}} \\
&= \|A - P\|_{S,2}^2 + e_{\bar{p}}^\top\left(\Theta(I - L) + (I - L)^\top\Theta\right) e_{\bar{p}}.
\end{aligned}$$

For the same reasonings as above, $e_{\bar{p}}^{\top}\left(\Theta(I-L)+(I-L)^{\top}\Theta\right)e_{\bar{p}}\geq 0$, and the result follows. □

Generalizations of the polar decomposition when using *deterministic and positive definite* inner products have been discussed in detail in [22].

The columns of the matrix P are $S^{\top}S$-orthogonal but not orthogonal. Hence, a natural question is how close P is to a matrix with orthonormal columns. To this end, let $P=Q_P H_P$ be the polar decomposition of P, where Q_P has orthonormal columns and is the closest such matrix to P. The following bound was shown in [21, Lemma 5.1] for a general P of full column rank,

$$\frac{\|P^{\top}P-I\|_2}{\|P\|_2+1}\leq\|P-Q_P\|_2\leq\|P^{\top}P-I\|_2. \tag{15}$$

For P being $S^{\top}S$-orthogonal, by using Proposition 1 we thus have, with high probability,

$$\|P-Q_P\|_2\leq\|P^{\top}P-I\|_2\leq\frac{\varepsilon}{1-\varepsilon}. \tag{16}$$

Analogously, given a matrix T with orthonormal columns, we can estimate how close T is to a matrix with $S^{\top}S$-orthonormal columns.

Lemma 2 *Given a matrix T with orthonormal columns, let Q_T be the $S^{\top}S$-orthogonal factor of the $S^{\top}S$-polar decomposition of the matrix T. Then with high probability*

$$\|T-Q_T\|_{S,2}\leq\varepsilon. \tag{17}$$

Proof We need to generalize the upper bound in (15). To this end, we follow the proof in [21, Lemma 5.1]. Let $T=Q_T H$ be the $S^{\top}S$-polar decomposition of T, with H symmetric and positive definite. In particular, $(ST)^{\top}ST=H^2$. With explicit computation we can verify that

$$T^{\top}S^{\top}ST-I=(T-Q_T)^{\top}S^{\top}S(T+Q_T)=(T-Q_T)^{\top}S^{\top}SQ_T(H+I). \tag{18}$$

Since $H+I$ is nonsingular, we can write

$$(T-Q_T)^{\top}S^{\top}SQ_T=(T^{\top}S^{\top}ST-I)(H+I)^{-1},$$

so that

$$\|(T-Q_T)^{\top}S^{\top}SQ_T\|_2\leq\|T^{\top}S^{\top}ST-I\|_2\|(H+I)^{-1}\|_2\leq\|T^{\top}S^{\top}ST-I\|_2.$$

Moreover,

$$\|T-Q_T\|_{S,2}=\|H-I\|_2=\|Q_T^{\top}S^{\top}SQ_T(H-I)\|_2=\|Q_T^{\top}S^{\top}S(T-Q_T)\|_2,$$

so that

$$\|T - Q_T\|_{S,2} \leq \|T^\top S^\top ST - I\|_2,$$

and the result follows from Proposition 2. □

We are now left to evaluate the quality of the minimizer when either orthogonality or sketched orthogonality is used.

Proposition 3 *Let P solve (11), and let T solve (1) in the 2-norm. With the previous notation, with high probability it holds*

$$\|A - T\|_2 - \frac{\varepsilon}{1-\varepsilon} \leq \|A - P\|_2 \leq \frac{1+\varepsilon}{1-\varepsilon}\|A - T\|_2 + \frac{\varepsilon}{1-\varepsilon}.$$

Proof Let Q_P be the orthogonal factor of the polar decomposition of P. Then

$$\|A - T\|_2 \leq \|A - Q_P\|_2 \leq \|A - P\|_2 + \|P - Q_P\|_2 \leq \|A - P\|_2 + \frac{\varepsilon}{1-\varepsilon},$$

where, in the last inequality, (16) was used. This gives the left-hand side bound. To obtain the bound on the right side, let Q_T be the $S^\top S$-orthogonal factor of the $S^\top S$-polar decomposition of T. Then

$$\|A - P\|_{S,2} \leq \|A - Q_T\|_{S,2} \leq \|A - T\|_{S,2} + \|T - Q_T\|_{S,2} \leq (1+\varepsilon)\|A - T\|_2 + \varepsilon$$

where in the last inequality, (17) was used.

The result follows from recalling that $(1-\varepsilon)\|A - P\|_2 \leq \|A - P\|_{S,2}$. □

Proposition 3 shows that when solving (11), the minimizer P attains an error in the 2-norm which is not too far from the best attainable error (in that norm). Moreover, the cheaper computation of the $S^T S$-SVD could also make the exact solution of (11) affordable in the case of large dimensional problems. Notice that from the proof of Proposition 3, it is evident that the factor $\frac{1+\varepsilon}{1-\varepsilon}$ arises in relating the matrix 2-norm and the S, 2-norm.

Example 3 We consider the matrix $A \in \mathbb{R}^{m\times n}$ stemming from the benchmark problem `abtaha2` in the SuiteSparse Matrix Collection Repository[7] The matrix has size 37932×331. We compute the nearest orthogonal matrix T to A in the 2-norm, that is, we solve (1) for $* = 2$, and compare the obtained solution with P coming from solving (11) for the S, 2 norm. To this end, we consider the sketching (10) for different values of s.

To construct P, we first compute the factors of the $S^\top S$-SVD of A, i.e., $A = W\Theta V^\top$ as shown in sect. 3. We then set $P = WV^\top$. Similarly, to obtain T, we compute the SVD of A, $A = U\Sigma Y^\top$, by the matlab function `svd`, and set $T = UY^\top$. This choice of T yields $\|A - T\|_2 = 24.77$, and the time devoted to its computation is 1.09 seconds.

In Table 2 we report the distance from A attained by P, measured in the 2-norm. This is not far from the one provided by T, as predicted by Proposition 3. Indeed, for

[7] Available at https://sparse.tamu.edu/.

Table 2 Example 3. 2-norm distance with the solution P to (11), 2-norm difference between P and the solution T to (1), and the run times for computing P, as the sketching dimension s varies. The reported results are averaged over 50 runs

s	$S^\top S$-orth $\|A-P\|_2$	Distance $\|P-T\|_2$	Time (s)
$2n$	24.99	3.69	0.46
$4n$	24.80	2.57	0.46
$6n$	24.77	2.36	0.47
$8n$	24.76	2.27	0.48
$10n$	24.75	2.21	0.48
$12n$	24.74	2.18	0.49

$\varepsilon = 0.5$, Proposition 3 states that, with high probability,

$$\|A-T\|_2 - 1 \leq \|A-P\|_2 \leq 3\|A-T\|_2 + 1,$$

and these bounds are satisfied for all tested values of s. We also report $\|P-T\|_2$ and the running times for computing P while varying s. Looking at the results in Table 2, we can see that $\|P-T\|_2$ is always rather moderate, with up to 55% cuts in running time when computing P instead of T.

6 Conclusions

We have formally introduced the randomization-based $S^\top S$-SVD decomposition of a given tall matrix A, which holds with high probability. This decomposition resembles the standard SVD, where, however, the left singular vectors are constrained to be orthonormal with respect to the $S^\top S$-norm. The $S^\top S$-SVD has then been employed to derive a number of results scattered in the literature, and to directly derive probabilistic bounds on the distance from (standard) orthogonality of the sketched orthogonal factor, also in state-of-the-art randomized algorithms for the QR factorization. We believe that the $S^\top S$-SVD has the potential to fully characterize the behavior of sketching techniques applied to least squares problems.

We have also studied the related problem of finding the nearest orthogonal matrix to A, in the same $S^\top S$-norm, and adopted the $S^\top S$-SVD of A for its solution. Additional comparison bounds have complemented our presentation, illustrating that sketched orthogonality allows one to obtain results comparable to those with Euclidean orthogonality, at a significantly lower computational cost.

Acknowledgements Both authors are members of the INdAM Research Group GNCS. Moreover, their work was partially supported by the European Union - NextGenerationEU under the National Recovery and Resilience Plan (PNRR) - Mission 4 Education and research - Component 2 From research to business - Investment 1.1 Notice Prin 2022 - DD N. 104 of 2/2/2022, entitled “Low-rank Structures and Numerical Methods in Matrix and Tensor Computations and their Application”, code 20227PCCKZ – CUP J53D23003620006.

Funding Open access funding provided by Alma Mater Studiorum - Università di Bologna within the CRUI-CARE Agreement.

Declarations

Conflicts of Interest Not applicable.

References

1. Absil, P.-A., Mahony, R., Sepulchre, R.: Optimization Algorithms on Matrix Manifolds. Princeton University Press, Princeton, NJ (2008)
2. Avron, H., Maymounkov, P., Toledo, S.: Blendenpik: supercharging LAPACK's least-squares solver. SIAM J. Sci. Comput. **32**, 1217–1236 (2010). https://doi.org/10.1137/090767911
3. Balabanov, O., Grigori, L.: Randomized gram-schmidt process with application to GMRES. SIAM J. Sci. Comput. **44**, A1450–A1474 (2022). https://doi.org/10.1137/20M138870X
4. Balabanov, O., Nouy, A.: Randomized linear algebra for model reduction, Part I: galerkin methods and error estimation. Adv. Comput. Math. **45**, 2969–3019 (2019). https://doi.org/10.1007/s10444-019-09725-6
5. Beckermann, B., Townsend, A.: On the singular values of matrices with displacement structure. SIAM J. Matrix Anal. Appl. **38**, 1227–1248 (2017). https://doi.org/10.1137/16M1096426
6. Björck, A.: Numerical methods for least squares problems. Soc. Industrial Appl. Math. (1996). https://doi.org/10.1137/1.9781611971484
7. Björck, A., Bowie, C.: An iterative algorithm for computing the best estimate of an orthogonal matrix. SIAM J. Numer. Anal. **8**, 358–364 (1971). https://doi.org/10.1137/0708036
8. Boutsidis, C., Woodruff, D.P.: Optimal CUR matrix decompositions. SIAM J. Comput. **46**, 543–589 (2017). https://doi.org/10.1137/140977898
9. Burke, L., Güttel, S., Soodhalter, K.M.: GMRES with randomized sketching and deflated restarting. SIAM J. Matrix Anal. Appl. **46**, 702–725 (2025). https://doi.org/10.1137/23M1619472
10. Celledoni, E., Owren, B.: A class of intrinsic schemes for orthogonal integration. SIAM J. Numer. Anal. **40**, 2069–2084 (2002)
11. Ceruti, G., Kusch, J., Lubich, C.: A rank-adaptive robust integrator for dynamical lowrank approximation. BIT Numer. Math. **62**, 1149–1174 (2022)
12. Edelman, A., Arias, T., Smith, S.T.: The geometry of algorithms with orthogonality constraints. SIAM J. Matrix Anal. Appl. **20**, 303–353 (1999)
13. Gilbert, A.C., Park, J.Y., Wakin, M.B.: Sketched SVD: Recovering Spectral Features from Compressive Measurements, in Proc. of SPARS 2013, 2013. (also arXiv:1211.0361, (2012))
14. Golub, G., Van Loan, C.F.: Matrix Computations, 4th edn. The Johns Hopkins University Press, Baltimore (2013)
15. Grigori, L., Timsit, E.: Randomized Householder QR, Tech. Report arXiv:2405.10923, (2024)
16. Grigori, L., Xue, Z.: Randomized strong rank-revealing QR for column subset selection and low-rank matrix approximation, Tech. Report arXiv:2503.18496, (2025)
17. Hairer, E., Lubich, C., Wanner, G.: Geometric numerical integration. Structure-preserving algorithms for ordinary differential equations, vol. 31 of Springer Series in Computational Mathematics, Springer, Berlin, (2002)
18. Halko, N., Martinsson, P.G., Tropp, J.A.: Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions. SIAM Rev. **53**, 217–288 (2011). https://doi.org/10.1137/090771806
19. Hansen, P.C.: Discrete inverse problems: insight and algorithms, SIAM, (2010)

20. Higgins, A.J., Szyld, D.B., Boman, E.G., Yamazaki, I.: Analysis of randomized Householder-Cholesky QR factorization with multisketching. Numer. Math. **157**, 1695–1737 (2025). https://doi.org/10.1007/00211-025-01492-5
21. Higham, N.J.: The matrix sign decomposition and its relation to the polar decomposition. Linear Algebra Appl. **212**(213), 3–20 (1994)
22. Higham, N.J., Mehl, C., Tisseur, F.: The canonical generalized polar decomposition. SIAM J. Matrix Anal. Appl. **31**, 2163–2180 (2010). https://doi.org/10.1137/090765018
23. Jang, Y., Grigori, L.: Randomized orthogonalization process with reorthogonalization. Num. Linear Algebra Appl. **32**, e70029 (2025). https://doi.org/10.1002/nla.70029
24. Kireeva, A., Tropp, J.A.: Randomized matrix computations: themes and variations, tech. report, arXiv:2402.17873, v.3, (2025), https://doi.org/10.7907/7YADE-5K351, https://authors.library.caltech.edu/records/4v7j7-msj10
25. Martinsson, P.-G., Tropp, J.A.: Randomized numerical linear algebra: foundations and algorithms. Acta Numer **29**, 403–572 (2020). https://doi.org/10.1017/S0962492920000021
26. Rokhlin, V., Tygert, M.: A fast randomized algorithm for overdetermined linear least-squares regression. Proc. Natl. Acad. Sci. **105**, 13212–13217 (2008). https://doi.org/10.1073/pnas.0804869105
27. Saibaba, A.K.: Randomized subspace iteration: analysis of canonical angles and unitarily invariant norms. SIAM J. Matrix Anal. Appl. **40**, 23–48 (2019). https://doi.org/10.1137/18M1179432
28. Sarlos, T.: Improved Approximation Algorithms for Large Matrices via Random Projections, in 2006 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS'06), pp. 143–152, (2006) https://doi.org/10.1109/FOCS.2006.37
29. Tropp, J.A.: Improved analysis of the subsampled randomized Hadamard transform. Adv. Adapt. Data Anal. **03**, 115–126 (2011). https://doi.org/10.1142/S1793536911000787
30. Van Loan, C.F.: Generalizing the singular value decomposition. SIAM J. Numer. Anal. **13**, 76–83 (1976)
31. Varah, J.M.: On the numerical solution of ill-conditioned linear systems with applications to ill-posed problems. SIAM J. Numer. Anal. **10**, 257–267 (1973)
32. Woolfe, F., Liberty, E., Rokhlin, V., Tygert, M.: A fast randomized algorithm for the approximation of matrices. Appl. Comput. Harmon. Anal. **25**, 335–366 (2008). https://doi.org/10.1016/j.acha.2007.12.002