

NEAR-OPTIMAL COLUMN-BASED MATRIX RECONSTRUCTION*

CHRISTOS BOUTSIDIS[†], PETROS DRINEAS[‡], AND MALIK MAGDON-ISMAIL[‡]

Abstract. We consider low-rank reconstruction of a matrix using a subset of its columns and present asymptotically optimal algorithms for both spectral norm and Frobenius norm reconstruction. The main tools we introduce to obtain our results are (i) the use of fast approximate SVD-like decompositions for column-based matrix reconstruction, and (ii) two deterministic algorithms for selecting rows from matrices with orthonormal columns, building upon the sparse representation theorem for decompositions of the identity that appeared in [J. D. Batson, D. A. Spielman, and N. Srivastava, *Twice-Ramanujan sparsifiers*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), 2009, pp. 255–262].

Key words. randomized algorithms, numerical linear algebra, low-rank approximations

AMS subject classifications. 15B52, 15A18, 11K45

DOI. 10.1137/12086755X

1. Introduction. The best rank- k approximation to a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is

$$\mathbf{A}_k = \sum_{i=1}^k \sigma_i \mathbf{u}_i \mathbf{v}_i^T,$$

where $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_k \geq 0$ are the top k singular values of \mathbf{A} , with associated left and right singular vectors $\mathbf{u}_i \in \mathbb{R}^m$ and $\mathbf{v}_i \in \mathbb{R}^n$, respectively. (See section 1.1 for notation and recall that the singular values and singular vectors of \mathbf{A} can be computed via the singular value decomposition (SVD) of \mathbf{A} in $O(mn \min\{m, n\})$ time.) It is well known that \mathbf{A}_k optimally approximates \mathbf{A} among all rank- k matrices, with respect to any unitarily invariant norm. There is considerable interest (e.g., [4, 6, 9, 7, 11, 15, 19, 20, 21]) in determining a minimum set of $r \ll n$ columns of \mathbf{A} which is approximately as good as \mathbf{A}_k at reconstructing \mathbf{A} . Such columns are important for interpreting data [21], building robust machine learning algorithms [4], feature selection, etc.

Let $\mathbf{A} \in \mathbb{R}^{m \times n}$ and let $\mathbf{C} \in \mathbb{R}^{m \times r}$ consist of r columns of \mathbf{A} for some $k \leq r < n$. We are interested in the reconstruction errors

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+ \mathbf{A}\|_\xi \quad \text{and} \quad \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi$$

for $\xi = 2, F$ (see section 1.1 for notation). The former is the reconstruction error for \mathbf{A} using the columns in \mathbf{C} ; the latter is the error from the best rank- k reconstruction of \mathbf{A} (under the appropriate norm) within the column space of \mathbf{C} . For fixed \mathbf{A} , k ,

*Received by the editors February 27, 2012; accepted for publication (in revised form) August 2, 2013; published electronically April 29, 2014.

<http://www.siam.org/journals/sicomp/43-2/86755.html>

[†]Mathematical Sciences Department, IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 (cboutsi@us.ibm.com). The work of this author was supported by the XDATA program of the Defense Advanced Research Projects Agency (DARPA), administered through Air Force Research Laboratory contract FA8750-12-C-0323.

[‡]Computer Science Department, Rensselaer Polytechnic Institute, Troy, NY 12180 (drinep@cs.rpi.edu, magdon@cs.rpi.edu). The work of this author was supported by NSF CCF-1016501 and NSF DMS-1008983.

and r , we would like these errors to be as close to

$$\|\mathbf{A} - \mathbf{A}_k\|_\xi$$

as possible. We present polynomial-time near-optimal constructions for arbitrary $r > k$, settling important open questions regarding column-based matrix reconstruction.

- *Spectral norm:* What is the best reconstruction error with $r > k$ columns? We present polynomial-time (deterministic and randomized) algorithms with approximation error asymptotically matching a lower bound proven in this work. Prior work had focused on the $r = k$ case and presented near-optimal polynomial-time algorithms [6, 17].
- *Frobenius norm:* How many columns are needed for relative error approximation, i.e., a reconstruction error of at most

$$(1 + \epsilon) \|\mathbf{A} - \mathbf{A}_k\|_F$$

for $\epsilon > 0$? We show that $O(k/\epsilon)$ columns contain a rank- k subspace which reconstructs \mathbf{A} to relative error, and we present the first sub-SVD (in terms of running time) randomized algorithm to identify these columns. This matches the $\Omega(k/\epsilon)$ lower bound in [9] and improves the best known upper bound of $O(k \log k + k/\epsilon)$ [6, 9, 12, 23].

1.1. Notation. $\mathbf{A}, \mathbf{B}, \dots$ are matrices; $\mathbf{a}, \mathbf{b}, \dots$ are column vectors. \mathbf{I}_n is the $n \times n$ identity matrix; $\mathbf{0}_{m \times n}$ is the $m \times n$ matrix of zeros; $\mathbf{1}_n$ is the $n \times 1$ vector of ones; \mathbf{e}_i is the standard basis (whose dimensionality will be clear from the context); $\text{rank}(\mathbf{A})$ is the rank of \mathbf{A} . The Frobenius and the spectral matrix norms are $\|\mathbf{A}\|_F^2 = \sum_{i,j} \mathbf{A}_{ij}^2$ and $\|\mathbf{A}\|_2 = \max_{\|\mathbf{x}\|_2=1} \|\mathbf{A}\mathbf{x}\|_2$; $\|\mathbf{A}\|_\xi$ is used if a result holds for both norms $\xi = 2$ and $\xi = F$. The SVD of \mathbf{A} , with $\text{rank}(\mathbf{A}) = \rho$, is

$$\mathbf{A} = \underbrace{\begin{pmatrix} \mathbf{U}_k & \mathbf{U}_{\rho-k} \end{pmatrix}}_{\mathbf{U}_A \in \mathbb{R}^{m \times \rho}} \underbrace{\begin{pmatrix} \Sigma_k & \mathbf{0} \\ \mathbf{0} & \Sigma_{\rho-k} \end{pmatrix}}_{\Sigma_A \in \mathbb{R}^{\rho \times \rho}} \underbrace{\begin{pmatrix} \mathbf{V}_k^T \\ \mathbf{V}_{\rho-k}^T \end{pmatrix}}_{\mathbf{V}_A^T \in \mathbb{R}^{\rho \times n}},$$

with singular values $\sigma_1 \geq \dots \geq \sigma_k \geq \sigma_{k+1} \geq \dots \geq \sigma_\rho > 0$. We will use $\sigma_i(\mathbf{A})$ to denote the i th singular value of \mathbf{A} when the matrix is not clear from the context. The matrices $\mathbf{U}_k \in \mathbb{R}^{m \times k}$ and $\mathbf{U}_{\rho-k} \in \mathbb{R}^{m \times (\rho-k)}$ contain the left singular vectors of \mathbf{A} , and, similarly, the matrices $\mathbf{V}_k \in \mathbb{R}^{n \times k}$ and $\mathbf{V}_{\rho-k} \in \mathbb{R}^{n \times (\rho-k)}$ contain the right singular vectors of \mathbf{A} . It is well known that $\mathbf{A}_k = \mathbf{U}_k \Sigma_k \mathbf{V}_k^T$ minimizes $\|\mathbf{A} - \mathbf{X}\|_\xi$ over all matrices $\mathbf{X} \in \mathbb{R}^{m \times n}$ of rank at most k . We use $\mathbf{A}_{\rho-k}$ to denote the matrix $\mathbf{A} - \mathbf{A}_k = \mathbf{U}_{\rho-k} \Sigma_{\rho-k} \mathbf{V}_{\rho-k}^T$. Also, $\mathbf{A}^+ = \mathbf{V}_A \Sigma_A^{-1} \mathbf{U}_A^T$ denotes the Moore–Penrose pseudoinverse of \mathbf{A} . For a symmetric positive definite matrix $\mathbf{A} = \mathbf{B}\mathbf{B}^T$, $\lambda_i(\mathbf{A}) = \sigma_i^2(\mathbf{B})$ denotes the i th eigenvalue of \mathbf{A} .

Finally, given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ and a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ with $r > k$, we formally define the matrix $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) \in \mathbb{R}^{m \times n}$ as the best approximation to \mathbf{A} within the column space of \mathbf{C} that has rank at most k . $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ minimizes the residual $\|\mathbf{A} - \hat{\mathbf{A}}\|_\xi$ over all $\hat{\mathbf{A}}$ in the column space of \mathbf{C} that have rank at most k (one can write $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) = \mathbf{C}\mathbf{X}$, where $\mathbf{X} \in \mathbb{R}^{r \times n}$ has rank at most k). In general,

$$\Pi_{\mathbf{C},k}^2(\mathbf{A}) \neq \Pi_{\mathbf{C},k}^F(\mathbf{A});$$

section 2.2 discusses the computation of $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$.

TABLE 1

Upper bounds for the approximation ratio $\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 / \|\mathbf{A} - \mathbf{A}_k\|_\xi^2$ for any $r > k$. Here $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ .

	Spectral norm ($\xi = 2$)	Frobenius norm ($\xi = \mathbb{F}$)
Deterministic	$\left(1 + \frac{1 + \sqrt{(\rho-k)/r}}{1 - \sqrt{k/r}}\right)^2$ (Thm. 1)	$1 + (1 - \sqrt{k/r})^{-2}$ (Thm. 2)
Randomized*	$O(n/r)$ (Thm. 3)	$1 + \frac{2k}{r}(1 + o(1))$ (Thm. 5)

*We give a bound on the expected approximation ratio.

1.2. Main results. Since

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_\xi \leq \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi,$$

we will state all our bounds in terms of the latter quantity. Note that we chose to state our Frobenius norm bounds in terms of the square of the Frobenius norm; this choice facilitates comparisons with prior work and simplifies our proofs (see also Table 1 for a summary of our results).

THEOREM 1 (deterministic spectral norm reconstruction). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ and a target rank $k < \rho$, there exists a deterministic polynomial-time algorithm to select $r > k$ columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2 \leq \left(1 + \frac{1 + \sqrt{(\rho-k)/r}}{1 - \sqrt{k/r}}\right) \|\mathbf{A} - \mathbf{A}_k\|_2 = O\left(\sqrt{\rho/r}\right) \|\mathbf{A} - \mathbf{A}_k\|_2.$$

The matrix \mathbf{C} can be computed in $T_{SVD} + O(rn(k^2 + (\rho - k)^2))$ time, where T_{SVD} is the time needed to compute all ρ right singular vectors of \mathbf{A} .

Our algorithm uses the matrices \mathbf{V}_k and $\mathbf{V}_{\rho-k}$ of the right singular vectors of \mathbf{A} . These matrices can be computed in $O(mn \min\{m, n\})$ time via the SVD. The asymptotic multiplicative error of the above theorem matches a lower bound that we prove in section 9.1. This is the first spectral reconstruction algorithm with asymptotically optimal guarantees for arbitrary $r > k$. Previous work presented near-optimal algorithms for $r = k$ [17]. We note that in section 4 we will present a result that achieves a slightly worse error bound (essentially replacing ρ by n in the accuracy guarantee), but only uses the top k right singular vectors of \mathbf{A} (i.e., the matrix \mathbf{V}_k from the SVD of \mathbf{A}).

THEOREM 2 (deterministic Frobenius norm reconstruction). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ and a target rank $k < \rho$, there exists a deterministic polynomial-time algorithm to select $r > k$ columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^\mathbb{F}(\mathbf{A})\|_\mathbb{F}^2 \leq \left(1 + (1 - \sqrt{k/r})^{-2}\right) \|\mathbf{A} - \mathbf{A}_k\|_\mathbb{F}^2.$$

The matrix \mathbf{C} can be computed in $T_{\mathbf{V}_k} + O(mn + nrk^2)$ time, where $T_{\mathbf{V}_k}$ is the time needed to compute the top k right singular vectors of \mathbf{A} .

Our bound implies a constant-factor approximation. Previous work presents deterministic near-optimal algorithms for $r = k$ [6]; we are unaware of any deterministic algorithms for $r > k$.

The next two theorems guarantee (up to small constant factors) the same bounds as Theorems 1 and 2, but the proposed algorithms are considerably more efficient.

In particular, there is no need to exactly compute the right singular vectors of \mathbf{A} , because approximations suffice.

THEOREM 3 (fast spectral norm reconstruction). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , a target rank $2 \leq k < \rho$, and $0 < \epsilon < 1$, there exists a randomized algorithm to select $r > k$ columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2] \leq (\sqrt{2} + \epsilon) \left(\frac{1 + \sqrt{n/r}}{1 - \sqrt{k/r}} \right) \|\mathbf{A} - \mathbf{A}_k\|_2 = O\left(\sqrt{n/r}\right) \|\mathbf{A} - \mathbf{A}_k\|_2.$$

The matrix \mathbf{C} can be computed in $O(mnk\epsilon^{-1} \log(k^{-1} \min\{m, n\}) + nrk^2)$ time.

THEOREM 4 (fast Frobenius norm reconstruction). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , a target rank $2 \leq k < \rho$, and $0 < \epsilon < 1$, there exists a randomized algorithm to select $r > k$ columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2] \leq (1 + \epsilon) \left(1 + \left(1 - \sqrt{k/r}\right)^{-2} \right) \|\mathbf{A} - \mathbf{A}_k\|_F^2.$$

The matrix \mathbf{C} can be computed in $O(mnk\epsilon^{-1} + nrk^2)$ time.

Our last yet perhaps most interesting result guarantees relative-error Frobenius norm approximation by combining the algorithm of Theorem 4 with one round of adaptive sampling [9, 7]. This is the *first* relative-error approximation for Frobenius norm reconstruction that uses a linear number of columns in k (the target rank). Previous work [12, 23, 9, 6] achieves relative error with $O(k \log k + k/\epsilon)$ columns. Our result asymptotically matches the $\Omega(k/\epsilon)$ lower bound in [9].

Notice that in Theorems 2 and 4, which use the deterministic spectral sparsification technique of Lemma 14 to select the columns, we only achieve a $2 + \epsilon$ error by selecting $O(k/\epsilon^2)$ columns. To improve this constant factor approximation to a relative error bound we used the adaptive sampling idea from [9, 7].

THEOREM 5 (fast relative-error Frobenius norm reconstruction). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , a target rank $2 \leq k < \rho$, and $0 < \epsilon < 1$, there exists a randomized algorithm to select at most*

$$r = \frac{2k}{\epsilon} (1 + o(1))$$

columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2] \leq (1 + \epsilon) \|\mathbf{A} - \mathbf{A}_k\|_F^2.$$

The matrix \mathbf{C} can be computed in $O((mnk + nk^3) \epsilon^{-2/3})$ time.

1.3. Running times. Notice that the running times in the theorems presented above are stated in terms of the number of operations needed to compute the matrix \mathbf{C} , and, for simplicity, we assume that \mathbf{A} is dense; if \mathbf{A} is sparse, additional savings might be possible. Our accuracy guarantees are in terms of the optimal matrix $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$, which would require additional time to compute. For the Frobenius norm, computing $\Pi_{\mathbf{C},k}^F(\mathbf{A})$ is straightforward and requires only an additional $O(mnr + (m+n)r^2)$ time (see the discussion in section 2.2). For the spectral norm, we are not aware of any efficient algorithm to compute $\Pi_{\mathbf{C},k}^2(\mathbf{A})$ exactly. In section 2.2 we present a simple approach that computes $\hat{\Pi}_{\mathbf{C},k}^2(\mathbf{A})$, a constant factor approximation to $\Pi_{\mathbf{C},k}^2(\mathbf{A})$, in $O(mnr + (m+n)r^2)$ time. Our bounds in Theorems 1 and 3 can be restated in terms of the error

$$\|\mathbf{A} - \hat{\Pi}_{\mathbf{C},k}^2(\mathbf{A})\|_2;$$

the accuracy guarantees only weaken by small constant factors.

TABLE 2
 Lower bounds for the approximation ratio $\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 / \|\mathbf{A} - \mathbf{A}_k\|_\xi^2$. Here $\mathbf{A} \in \mathbb{R}^{m \times n}$.

# Columns (r)	Spectral norm ($\xi = 2$)	Frobenius norm ($\xi = F$)
$r = k$	n/k [6]	$k + 1$ [7]
$r > k$	n/r (section 9.1)	$1 + k/r$ [9] (and section 9.2)

1.4. Lower bounds. Table 2 provides a summary of lower bounds for the ratio

$$\frac{\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2}{\|\mathbf{A} - \mathbf{A}_k\|_\xi^2},$$

where \mathbf{C} is a matrix consisting of r columns of \mathbf{A} , with $r \geq k$. Theorem 35 contributes a new lower bound for the spectral norm case when $r > k$. Note that any lower bound for the ratio $\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_\xi^2 / \|\mathbf{A} - \mathbf{A}_k\|_\xi^2$ implies the same lower bound for $\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 / \|\mathbf{A} - \mathbf{A}_k\|_\xi^2$; the converse, however, is not true.

1.5. Prior results on column-based matrix reconstructions. There is a long literature on algorithms for column-based matrix reconstruction using $r \geq k$ columns. The first result goes back to [16], with the most recent one being, to the best of our knowledge, the work in [6].

1.5.1. The Frobenius norm case. We present known upper bounds for the approximation ratio

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 / \|\mathbf{A} - \mathbf{A}_k\|_F^2.$$

We start with the $r = k$ case. In [3] a $T_{\mathbf{V}_k} + O(nk + k^3(\log^2 k)(\log \log k))$ time randomized algorithm is described which provides an upper bound $O(k \log^{\frac{1}{2}} k)$ with constant probability. This bound was subsequently improved in [6]. More precisely, Theorem 8 of [6] gives a $(k+1)$ deterministic approximation running in $O(knm^3 \log m)$ time; this upper bound matches the lower bound in [7]. Also presented in [6] are three randomized algorithms such that $\mathbb{E}[\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2] = (k+1)\|\mathbf{A} - \mathbf{A}_k\|_F^2$. These randomized algorithms are presented in Theorem 7, Proposition 16, and Proposition 18 and run in $O(knm^3 \log m)$, $O(kn^3m + kn^4 \log n)$, and $O(kT_{SVD} + knm^2)$ time, respectively. Moreover, Theorem 9 in [6] presents a randomized algorithm that runs in time $O(mn \log nk^2 \epsilon^{-2} + n \log^3 n \cdot k^7 \epsilon^{-6} \log(k\epsilon^{-1} \log n))$ such that with constant probability $\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 \leq (1+\epsilon) \cdot (k+1)\|\mathbf{A} - \mathbf{A}_k\|_F^2$ for any $0 < \epsilon < 1$. Finally, [18] improved upon the running time of the results in [6]. More precisely, Theorem 2 in [18] gives an $O(knm^2)$ time randomized algorithm with a $(k+1)$ multiplicative error (in expectation).

When $r = \Omega(k \log k)$, relative-error approximations are known. In [12] the first result that achieved such a bound was presented, using random sampling of the columns of \mathbf{A} according to the Euclidean norms of the rows of \mathbf{V}_k , the so-called leverage scores [21]. More specifically, a $(1 + \epsilon)$ approximation was proven using $r = \Omega(k\epsilon^{-2} \log(k\epsilon^{-1}))$ columns in $T_{\mathbf{V}_k} + O(kn + r \log r)$ time. The paper [23] argued that the same technique gives a $(1 + \epsilon)$ approximation using $r = \Omega(k \log k + k\epsilon^{-1})$ columns. It also showed how to improve the running time to $T_{\tilde{\mathbf{V}}_k} + O(kn + r \log r)$, where $\tilde{\mathbf{V}}_k \in \mathbb{R}^{n \times k}$ contains the right singular vectors of an approximation to \mathbf{A}_k and

can be computed in $o(mn \min\{m, n\})$ time, which is less than the time needed to compute the SVD of \mathbf{A} . In [9], the authors leveraged volume sampling and presented an approach that achieves a relative-error approximation using $O(k^2 \log k + k\epsilon^{-1})$ columns in $O(mnk^2 \log k)$ time. Also, it is possible to combine the fast volume sampling approach in [6] (setting, for example, $\epsilon = 1/2$) with $O(\log k)$ rounds of adaptive sampling as described in [9] to achieve a relative-error approximation using $O(k \log k + k\epsilon^{-1})$ columns. The running time of this combined algorithm is $O(mnk^2 \log n + nk^7 \log^3 n \cdot \log(k \log n))$. The techniques in [12] do not apply to general $r > k$, since $\Omega(k \log k)$ columns must be sampled in order to preserve rank with random sampling.

A related line of work (including [8, 13, 14, 24]) has focused on the construction of coresets and sketches for high-dimensional subspace approximation with respect to general ℓ_p norms. In our setting, $p = 2$ corresponds to Frobenius norm matrix reconstruction, and Theorem 1.3 of [24] presents an exponential in k/ϵ algorithm to select $O(k^2 \epsilon^{-1} \log(k/\epsilon))$ columns that guarantee a relative-error approximation. It would be interesting to understand whether the techniques of [8, 13, 14, 24] can be extended to match our results here in the special case of $p = 2$.

The recent work in [18] presents a deterministic and a randomized algorithm for arbitrary $r \geq k$ that guarantee upper bounds for the ratio $\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbb{F}}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2$. More precisely, Theorem 1 in [18] presents an $O(rnm^3 \log m)$ time deterministic algorithm with bound $(r + 1)/(r + 1 - k)$, which is tight up to low order terms if $r = o(n)$. Also, Theorem 2 in [18] presents an $O(rnm^2)$ time randomized algorithm which achieves the same bound in expectation. We should note that it is not obvious how to extend the results in [18] to obtain comparable bounds for the ratio $\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2$.

1.5.2. The spectral norm case. We present known guarantees for the approximation ratio

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2^2 / \|\mathbf{A} - \mathbf{A}_k\|_2^2.$$

In general, results for spectral norm have been rare. When $r = k$, the strongest bound emerges from the strong rank revealing QR (RRQR) [17] (specifically Algorithm 4 in [17]), which, for $f > 1$, runs in $O(mnk \log_f n)$ time and guarantees an $f^2 k(n - k) + 1$ approximation. For $r > k$, to the best of our knowledge, there is no easy way to extend the RRQR guarantees. In fact we are only aware of one bound (see below) that is applicable to this domain, other than those obtained by trivially extending the Frobenius norm bounds, because any α -approximation in the Frobenius norm gives an $\alpha(\rho - k)$ approximation in the spectral norm:

$$\begin{aligned} \|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2^2 &\leq \|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_2^2 \leq \|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2 \\ &\leq \alpha \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2 \leq \alpha(\rho - k) \|\mathbf{A} - \mathbf{A}_k\|_2^2. \end{aligned}$$

The recent work [1] describes a deterministic $T_{\mathbf{V}_k} + O(nk(n - r))$ time algorithm that guarantees an approximation error

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2^2 / \|\mathbf{A} - \mathbf{A}_k\|_2^2 \leq 2 + k(n - r)/(r - k + 1)$$

for any $r \geq k$.

2. Matrix norm properties and the computation of $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$.

2.1. Matrix norm properties. Recall notation from section 1.1; for any matrix \mathbf{A} of rank at most ρ , it is well known that $\|\mathbf{A}\|_F^2 = \sum_{i=1}^\rho \sigma_i^2(\mathbf{A})$ and $\|\mathbf{A}\|_2 = \sigma_1(\mathbf{A})$. Also, the best rank- k approximation to \mathbf{A} satisfies $\|\mathbf{A} - \mathbf{A}_k\|_2 = \sigma_{k+1}(\mathbf{A})$ and $\|\mathbf{A} - \mathbf{A}_k\|_F^2 = \sum_{i=k+1}^\rho \sigma_i^2(\mathbf{A})$. For any two matrices \mathbf{A} and \mathbf{B} of appropriate dimensions, $\|\mathbf{A}\|_2 \leq \|\mathbf{A}\|_F \leq \sqrt{\rho}\|\mathbf{A}\|_2$, $\|\mathbf{AB}\|_F \leq \|\mathbf{A}\|_F\|\mathbf{B}\|_2$, and $\|\mathbf{AB}\|_F \leq \|\mathbf{A}\|_2\|\mathbf{B}\|_F$. The latter two properties are stronger versions of the standard submultiplicativity property. We refer to the next lemma as matrix-Pythagoras.

LEMMA 6. *If $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{m \times n}$ and $\mathbf{XY}^T = \mathbf{0}_{m \times m}$ or $\mathbf{X}^T\mathbf{Y} = \mathbf{0}_{n \times n}$, then*

$$\begin{aligned} \|\mathbf{X} + \mathbf{Y}\|_F^2 &= \|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2, \\ \max\{\|\mathbf{X}\|_2^2, \|\mathbf{Y}\|_2^2\} &\leq \|\mathbf{X} + \mathbf{Y}\|_2^2 \leq \|\mathbf{X}\|_2^2 + \|\mathbf{Y}\|_2^2. \end{aligned}$$

Proof. Suppose $\mathbf{XY}^T = \mathbf{0}_{m \times m}$. Then $(\mathbf{X} + \mathbf{Y})(\mathbf{X} + \mathbf{Y})^T = \mathbf{XX}^T + \mathbf{YY}^T$. For $\xi = F$,

$$\|\mathbf{X} + \mathbf{Y}\|_F^2 = \text{Tr}((\mathbf{X} + \mathbf{Y})(\mathbf{X} + \mathbf{Y})^T) = \text{Tr}(\mathbf{XX}^T + \mathbf{YY}^T) = \|\mathbf{X}\|_F^2 + \|\mathbf{Y}\|_F^2.$$

For $\xi = 2$, let \mathbf{z} be any vector in \mathbb{R}^m . Then

$$\|\mathbf{X} + \mathbf{Y}\|_2^2 = \max_{\|\mathbf{z}\|_2=1} \mathbf{z}^T(\mathbf{X} + \mathbf{Y})(\mathbf{X} + \mathbf{Y})^T\mathbf{z} = \max_{\|\mathbf{z}\|_2=1} (\mathbf{z}^T\mathbf{XX}^T\mathbf{z} + \mathbf{z}^T\mathbf{YY}^T\mathbf{z}).$$

We have that $\max_{\|\mathbf{z}\|_2=1} (\mathbf{z}^T\mathbf{XX}^T\mathbf{z} + \mathbf{z}^T\mathbf{YY}^T\mathbf{z})$ is at most

$$\max_{\|\mathbf{z}\|_2=1} \mathbf{z}^T\mathbf{XX}^T\mathbf{z} + \max_{\|\mathbf{z}\|_2=1} \mathbf{z}^T\mathbf{YY}^T\mathbf{z} = \|\mathbf{X}\|_2^2 + \|\mathbf{Y}\|_2^2$$

and that

$$\max_{\|\mathbf{z}\|_2=1} (\mathbf{z}^T\mathbf{XX}^T\mathbf{z} + \mathbf{z}^T\mathbf{YY}^T\mathbf{z}) \geq \max_{\|\mathbf{z}\|_2=1} \mathbf{z}^T\mathbf{XX}^T\mathbf{z} = \|\mathbf{X}\|_2^2,$$

since $\mathbf{z}^T\mathbf{YY}^T\mathbf{z}$ is nonnegative for any vector \mathbf{z} . We get the same lower bound with $\|\mathbf{Y}\|_2^2$ instead, which means we can lower bound by $\max\{\|\mathbf{X}\|_2^2, \|\mathbf{Y}\|_2^2\}$. The case when $\mathbf{X}^T\mathbf{Y} = \mathbf{0}_{n \times n}$ can be proven similarly. \square

A projection operator \mathbf{P} equals its own square, $\mathbf{P}^2 = \mathbf{P}$. Projection operators play an important role in our analysis. The following lemma is well known for nontrivial symmetric projection matrices, but also holds for nonsymmetric (oblique) projection matrices.

LEMMA 7 (see [25]). *Let \mathbf{P} be a nonnull projection. Then $\|\mathbf{I} - \mathbf{P}\|_2 \leq \|\mathbf{P}\|_2$.*

(If in addition $\mathbf{I} - \mathbf{P}$, also a projection, is nonnull, then we get equality in the above lemma.)

2.2. Computing the best rank- k approximation $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$. Let $\mathbf{A} \in \mathbb{R}^{m \times n}$, let $k < n$ be an integer, and let $\mathbf{C} \in \mathbb{R}^{m \times r}$ with $r > k$. Recall that $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) \in \mathbb{R}^{m \times n}$ is the best rank- k approximation to \mathbf{A} in the column space of \mathbf{C} . We can write $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) = \mathbf{CX}^\xi$, where

$$\mathbf{X}^\xi = \underset{\Psi \in \mathbb{R}^{r \times n}: \text{rank}(\Psi) \leq k}{\text{argmin}} \|\mathbf{A} - \mathbf{C}\Psi\|_\xi^2.$$

In order to compute (or approximate) $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ given \mathbf{A} , \mathbf{C} , and k , we will use the following algorithm:

- 1: Orthonormalize the columns of \mathbf{C} in $O(mr^2)$ time to construct the matrix $\mathbf{Q} \in \mathbb{R}^{m \times r}$.
- 2: Compute $(\mathbf{Q}^T \mathbf{A})_k \in \mathbb{R}^{r \times n}$ via SVD in $O(mnr + nr^2)$ —the best rank- k approximation of $\mathbf{Q}^T \mathbf{A}$.
- 3: Return $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k \in \mathbb{R}^{m \times n}$ in $O(mnk)$ time.

Clearly, $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ is a rank- k matrix that lies in the column span of \mathbf{C} . Note that though $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ can depend on ξ , our algorithm computes the same matrix, independent of ξ . The next lemma, which is essentially Lemma 4.3 in [5] combined with an improvement of Theorem 9.3 in [19], proves that this algorithm computes $\Pi_{\mathbf{C},k}^F(\mathbf{A})$ and a constant factor approximation to $\Pi_{\mathbf{C},k}^2(\mathbf{A})$.

LEMMA 8. *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{C} \in \mathbb{R}^{m \times r}$, and an integer k , the matrix $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k \in \mathbb{R}^{m \times n}$ described above (where \mathbf{Q} is an orthonormal basis for the columns of \mathbf{C}) can be computed in $O(mnr + (m+n)r^2)$ time and satisfies*

$$\begin{aligned} \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_F^2 &= \|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2, \\ \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_2^2 &\leq 2\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2^2. \end{aligned}$$

Proof. Our proof for the Frobenius norm case is a mild modification of the proof of Lemma 4.3 in [5]. First, note that $\Pi_{\mathbf{C},k}^F(\mathbf{A}) = \Pi_{\mathbf{Q},k}^F(\mathbf{A})$, because $\mathbf{Q} \in \mathbb{R}^{m \times r}$ is an orthonormal basis for the column space of \mathbf{C} . Thus,

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 = \|\mathbf{A} - \Pi_{\mathbf{Q},k}^F(\mathbf{A})\|_F^2 = \min_{\Psi: \text{rank}(\Psi) \leq k} \|\mathbf{A} - \mathbf{Q}\Psi\|_F^2.$$

Now, using matrix-Pythagoras and the orthonormality of \mathbf{Q} ,

$$\|\mathbf{A} - \mathbf{Q}\Psi\|_F^2 = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^T \mathbf{A} + \mathbf{Q}(\mathbf{Q}^T \mathbf{A} - \Psi)\|_F^2 = \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^T \mathbf{A}\|_F^2 + \|\mathbf{Q}^T \mathbf{A} - \Psi\|_F^2.$$

Setting $\Psi = (\mathbf{Q}^T \mathbf{A})_k$ minimizes the above quantity over all rank- k matrices Ψ . Thus, combining the above results, $\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 = \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_F^2$.

We proceed to the spectral norm part of the proof, which combines ideas from Theorem 9.3 of [19] and matrix-Pythagoras to manipulate the term $\|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_2^2$. In the derivations below there are two sources of errors: the first comes from projecting \mathbf{A} on $\mathbf{Q}\mathbf{Q}^T$ and the second from taking a low-rank approximation of this projection:

$$\begin{aligned} \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_2^2 &= \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^T \mathbf{A} + \mathbf{Q}(\mathbf{Q}^T \mathbf{A} - (\mathbf{Q}^T \mathbf{A})_k)\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{Q}\mathbf{Q}^T \mathbf{A}\|_2^2 + \|\mathbf{Q}\mathbf{Q}^T \mathbf{A} - (\mathbf{Q}\mathbf{Q}^T \mathbf{A})_k\|_2^2 \\ &\stackrel{(a)}{\leq} \|\mathbf{A} - \Pi_{\mathbf{Q},k}^2(\mathbf{A})\|_2^2 + \|\mathbf{A} - \mathbf{A}_k\|_2^2 \\ &\leq 2\|\mathbf{A} - \Pi_{\mathbf{Q},k}^2(\mathbf{A})\|_2^2. \end{aligned}$$

The first inequality follows from the simple fact that $(\mathbf{Q}\mathbf{Q}^T \mathbf{A})_k = \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ and from matrix-Pythagoras; the first term in (a) follows because $\mathbf{Q}\mathbf{Q}^T \mathbf{A}$ is the (unconstrained, not necessarily of rank at most k) best approximation to \mathbf{A} in the column space of \mathbf{Q} ; the second term in (a) follows because $\mathbf{Q}\mathbf{Q}^T$ is a projector matrix and thus

$$\|\mathbf{Q}\mathbf{Q}^T \mathbf{A} - (\mathbf{Q}\mathbf{Q}^T \mathbf{A})_k\|_2^2 = \sigma_{k+1}^2(\mathbf{Q}\mathbf{Q}^T \mathbf{A}) \leq \sigma_{k+1}^2(\mathbf{A}) = \|\mathbf{A} - \mathbf{A}_k\|_2^2.$$

The last inequality follows because $\|\mathbf{A} - \mathbf{A}_k\|_2^2 \leq \|\mathbf{A} - \Pi_{\mathbf{Q},k}^2(\mathbf{A})\|_2^2$. \square

3. Main tools. Our two main tools are the use of matrix factorizations for column-based low-rank matrix reconstruction, and two deterministic sparsification lemmas which extend the work of [2].

3.1. Matrix factorizations. Our first tool (see Lemmas 9, 11, and 12) connects matrix factorizations and matrix reconstruction from its columns. Specifically, Lemmas 11 and 12 consider factorizations of the matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ of the form

$$\mathbf{A} = \mathbf{B}\mathbf{Z}^T + \mathbf{E},$$

where $\mathbf{B} = \mathbf{A}\mathbf{Z} \in \mathbb{R}^{m \times k}$, $\mathbf{Z} \in \mathbb{R}^{n \times k}$, $\mathbf{E} \in \mathbb{R}^{m \times n}$, and \mathbf{Z} has orthonormal columns. Note that the factorization decomposes \mathbf{A} to its projection $\mathbf{A}\mathbf{Z}\mathbf{Z}^T$ onto the k -dimensional space spanned by the columns of \mathbf{Z} and the orthogonal error $\mathbf{E} = \mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$, which implies that $\mathbf{E}\mathbf{Z} = \mathbf{0}_{m \times k}$. Lemma 9 shows how these factorizations are connected to column selection. Lemma 9 is the starting point of all our column reconstruction results.

LEMMA 9. *Let $\mathbf{A} = \mathbf{B}\mathbf{Z}^T + \mathbf{E}$, with $\mathbf{B} = \mathbf{A}\mathbf{Z}$ and $\mathbf{Z}^T\mathbf{Z} = \mathbf{I}_k$. Let $\mathbf{S} \in \mathbb{R}^{n \times r}$ be any matrix such that $\text{rank}(\mathbf{Z}^T\mathbf{S}) = \text{rank}(\mathbf{Z}) = k$. Let $\mathbf{C} = \mathbf{A}\mathbf{S} \in \mathbb{R}^{m \times r}$. Then*

$$(1) \quad \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{E}\|_\xi^2 + \|\mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\|_\xi^2$$

and

$$(2) \quad \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{E}\|_\xi^2 \cdot \|\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\|_2^2.$$

Proof. We first prove (1). The optimality of $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ implies that

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{A} - \mathbf{X}\|_\xi^2$$

over all matrices $\mathbf{X} \in \mathbb{R}^{m \times n}$ of rank at most k in the column space of \mathbf{C} . Consider the matrix $\mathbf{X} = \mathbf{C}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T$ (clearly \mathbf{X} is in the column space of \mathbf{C} and $\text{rank}(\mathbf{X}) \leq k$ because $\mathbf{Z} \in \mathbb{R}^{n \times k}$):

$$\begin{aligned} \|\mathbf{A} - \mathbf{C}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\|_\xi^2 &= \|\underbrace{\mathbf{B}\mathbf{Z}^T + \mathbf{E}}_{\mathbf{A}} - \underbrace{(\mathbf{B}\mathbf{Z}^T + \mathbf{E})\mathbf{S}}_{\mathbf{C} = \mathbf{A}\mathbf{S} = (\mathbf{B}\mathbf{Z}^T + \mathbf{E})\mathbf{S}} (\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\|_\xi^2 \\ &= \|\mathbf{B}\mathbf{Z}^T - \mathbf{B}\mathbf{Z}^T\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T + \mathbf{E} - \mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\|_\xi^2 \\ (3) \quad &\stackrel{(a)}{=} \|\mathbf{E} - \mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\|_\xi^2 \\ &\stackrel{(b)}{\leq} \|\mathbf{E}\|_\xi^2 + \|\mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\|_\xi^2. \end{aligned}$$

(a) follows because, by assumption, $\text{rank}(\mathbf{Z}^T\mathbf{S}) = k$, and thus $(\mathbf{Z}^T\mathbf{S})(\mathbf{Z}^T\mathbf{S})^+ = \mathbf{I}_k$, which implies that the first two terms cancel:

$$\mathbf{B}\mathbf{Z}^T - \mathbf{B}(\mathbf{Z}^T\mathbf{S})(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T = \mathbf{0}_{m \times n}.$$

(b) follows by matrix-Pythagoras because

$$\mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\mathbf{Z}^T\mathbf{E}^T = \mathbf{0}_{m \times n}$$

(recall that $\mathbf{E} = \mathbf{A}(\mathbf{I} - \mathbf{Z}\mathbf{Z}^T)$ and so $\mathbf{E}\mathbf{Z} = \mathbf{0}_{m \times k}$). The lemma follows by strong submultiplicativity because \mathbf{Z} has orthonormal columns, and hence $\|\mathbf{Z}\|_2 = 1$.

We now prove (2). In the above derivation up to (3), we have shown

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{A} - \mathbf{C}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top\|_\xi^2 = \|\mathbf{E}(\mathbf{I} - \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top)\|_\xi^2.$$

By strong submultiplicativity, the last term is bounded by

$$\|\mathbf{E}\|_\xi^2 \|\mathbf{I} - \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top\|_2^2 \leq \|\mathbf{E}\|_\xi^2 \|\mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top\|_2^2 = \|\mathbf{E}\|_\xi^2 \|\mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+\|_2^2.$$

The first step follows because

$$\mathbf{P} = \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top$$

is a nonnull projection because

$$\mathbf{P}^2 = \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top = \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ \mathbf{Z}^\top$$

(where we used $\mathbf{Z}^\top \mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+ = \mathbf{I}_k$), and so we can apply Lemma 7; the second step is because \mathbf{Z} is orthogonal. \square

In this work, we view \mathbf{C} as a dimensionally reduced or sampled sketch of \mathbf{A} ; \mathbf{S} is the dimension-reduction or sampling matrix. In other words, Lemma 9 argues that if the matrix \mathbf{S} preserves the rank of an approximate factorization of the original matrix \mathbf{A} , then the reconstruction of \mathbf{A} from $\mathbf{C} = \mathbf{A}\mathbf{S}$ has an error that is essentially proportional to the error of the approximate factorization. The importance of this lemma is that it indicates an algorithm for matrix reconstruction using a subset of the columns of \mathbf{A} . First, compute *any* factorization of the form $\mathbf{A} = \mathbf{B}\mathbf{Z}^\top + \mathbf{E}$, where $\mathbf{B} = \mathbf{A}\mathbf{Z}$ and $\|\mathbf{E}\|_\xi$ is small. Then compute a sampling matrix \mathbf{S} which satisfies the rank assumption and controls the error $\|\mathbf{E}\mathbf{S}(\mathbf{Z}^\top \mathbf{S})^+\|_\xi$.

An immediate corollary of Lemma 9 emerges by considering the SVD of \mathbf{A} . More specifically, consider the following factorization of \mathbf{A} :

$$\mathbf{A} = \mathbf{A}\mathbf{V}_k \mathbf{V}_k^\top + (\mathbf{A} - \mathbf{A}_k),$$

where \mathbf{V}_k is the matrix of the top k right singular vectors of \mathbf{A} . In the parlance of Lemma 9, $\mathbf{Z} = \mathbf{V}_k$, $\mathbf{B} = \mathbf{A}\mathbf{V}_k$, $\mathbf{E} = \mathbf{A} - \mathbf{A}_k$, and clearly $\mathbf{E}\mathbf{Z} = \mathbf{0}_{m \times k}$.

LEMMA 10. *Let $\mathbf{S} \in \mathbb{R}^{n \times r}$ be a matrix such that $\text{rank}(\mathbf{V}_k^\top \mathbf{S}) = k$. Let $\mathbf{C} = \mathbf{A}\mathbf{S}$. Then*

$$(4) \quad \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 + \|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}(\mathbf{V}_k^\top \mathbf{S})^+\|_\xi^2$$

and

$$(5) \quad \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_\xi^2 \leq \|\mathbf{A} - \mathbf{A}_k\|_\xi^2 \cdot \|\mathbf{S}(\mathbf{V}_k^\top \mathbf{S})^+\|_2^2.$$

The above lemma will be useful for designing the deterministic (spectral norm and Frobenius norm) column reconstruction algorithms of Theorems 1 and 2. However, computing the SVD is costly, and thus we would like to design a factorization of the form $\mathbf{A} = \mathbf{B}\mathbf{Z}^\top + \mathbf{E}$ that is as good as the SVD, but can be computed in $O(mnk)$ time. The next two lemmas achieve this goal by extending the algorithms in [19, 22] (see sections 5 and 6 for their proofs). We will use these factorizations to design fast column reconstruction algorithms in Theorems 3, 4, and 5.

LEMMA 11 (randomized fast spectral norm SVD). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , a target rank $2 \leq k < \rho$, and $0 < \epsilon < 1$, there exists an algorithm that computes a factorization $\mathbf{A} = \mathbf{B}\mathbf{Z}^\top + \mathbf{E}$, with $\mathbf{B} = \mathbf{A}\mathbf{Z}$, $\mathbf{Z}^\top \mathbf{Z} = \mathbf{I}_k$, and $\mathbf{E}\mathbf{Z} = \mathbf{0}_{m \times k}$ such that*

$$\mathbb{E}[\|\mathbf{E}\|_2] \leq (\sqrt{2} + \epsilon) \|\mathbf{A} - \mathbf{A}_k\|_2.$$

The proposed algorithm runs in $O(mnk\epsilon^{-1} \log(k^{-1} \min\{m, n\}))$ time.

LEMMA 12 (randomized fast Frobenius norm SVD). *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , a target rank $2 \leq k < \rho$, and $0 < \epsilon < 1$, there exists an algorithm that computes a factorization $\mathbf{A} = \mathbf{BZ}^T + \mathbf{E}$, with $\mathbf{B} = \mathbf{AZ}$, $\mathbf{Z}^T\mathbf{Z} = \mathbf{I}_k$, and $\mathbf{EZ} = \mathbf{0}_{m \times k}$ such that*

$$\mathbb{E} [\|\mathbf{E}\|_{\mathbb{F}}^2] \leq (1 + \epsilon)\|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2.$$

The proposed algorithm runs in $O(mnk\epsilon^{-1})$ time.

3.2. Sparse approximate decompositions of the identity. Lemmas 9, 11, and 12 argue that in order to achieve almost-optimal column-based matrix reconstruction, we need a sampling matrix \mathbf{S} that preserves the rank of \mathbf{Z} and controls the error $\|\mathbf{ES}(\mathbf{Z}^T\mathbf{S})^+\|_{\xi}$. We present algorithms for computing such a matrix \mathbf{S} in Lemmas 13 and 14. These lemmas were motivated by an important linear-algebraic result for a decomposition of the identity presented by Batson, Spielman, and Srivastava [2]. It is worth emphasizing that the result of [2] cannot be directly applied to the column reconstruction problem. Indeed, in our setting, it is necessary to control properties related to both matrices \mathbf{Z} and $\mathbf{E} = \mathbf{A} - \mathbf{BZ}^T$ simultaneously. In the spectral norm reconstruction case, we need to control the singular values of the two matrices; in the Frobenius norm reconstruction case, we need to control singular values and Frobenius norms of two matrices. The following two lemmas are proven in sections 7 and 8.

LEMMA 13 (dual set spectral sparsification). *Let $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$ be two equal cardinality decompositions of the identity, where $\mathbf{v}_i \in \mathbb{R}^k$ ($k < n$), $\mathbf{u}_i \in \mathbb{R}^\ell$ ($\ell \leq n$), $\sum_{i=1}^n \mathbf{v}_i\mathbf{v}_i^T = \mathbf{I}_k$, and $\sum_{i=1}^n \mathbf{u}_i\mathbf{u}_i^T = \mathbf{I}_\ell$. Given an integer r with $k < r \leq n$, there exists a set of weights $s_i \geq 0$ ($i = 1, \dots, n$), at most r of which are nonzero, such that*

$$\lambda_k \left(\sum_{i=1}^n s_i \mathbf{v}_i \mathbf{v}_i^T \right) \geq \left(1 - \sqrt{\frac{k}{r}} \right)^2 \quad \text{and} \quad \lambda_1 \left(\sum_{i=1}^n s_i \mathbf{u}_i \mathbf{u}_i^T \right) \leq \left(1 + \sqrt{\frac{\ell}{r}} \right)^2.$$

The weights s_i can be computed deterministically in $O(rn(k^2 + \ell^2))$ time.

In matrix notation, let \mathbf{U} and \mathbf{V} be the matrices whose rows are the vectors \mathbf{u}_i and \mathbf{v}_i , respectively. We can now construct the sampling matrix $\mathbf{S} \in \mathbb{R}^{n \times r}$ as follows: for $i = 1, \dots, n$, if s_i is nonzero, then include $\sqrt{s_i}\mathbf{e}_i$ as a column of \mathbf{S} ; here \mathbf{e}_i is the i th standard basis vector.¹ Using this matrix notation,

$$\sum_{i=1}^n s_i \mathbf{v}_i \mathbf{v}_i^T = \mathbf{V}^T \mathbf{S} \mathbf{S}^T \mathbf{V} \quad \text{and} \quad \sum_{i=1}^n s_i \mathbf{u}_i \mathbf{u}_i^T = \mathbf{U}^T \mathbf{S} \mathbf{S}^T \mathbf{U},$$

and so the above lemma guarantees that

$$\sigma_k(\mathbf{V}^T \mathbf{S}) \geq 1 - \sqrt{k/r} \quad \text{and} \quad \sigma_1(\mathbf{U}^T \mathbf{S}) \leq 1 + \sqrt{\ell/r}.$$

Clearly, \mathbf{S} may be viewed as a matrix that samples and rescales r rows of \mathbf{U} and \mathbf{V} (columns of \mathbf{U}^T and \mathbf{V}^T), namely, the rows that correspond to nonzero weights s_i .

LEMMA 14 (dual set spectral Frobenius sparsification.). *Let $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ be a decomposition of the identity, where $\mathbf{v}_i \in \mathbb{R}^k$ ($k < n$) and $\sum_{i=1}^n \mathbf{v}_i\mathbf{v}_i^T = \mathbf{I}_k$; let $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ be an arbitrary set of vectors, where $\mathbf{a}_i \in \mathbb{R}^\ell$. Then, given an integer*

¹Note that we slightly abused notation: indeed, the number of columns of \mathbf{S} is less than or equal to r , since at most r of the weights are nonzero. Here, we use r to also denote the actual number of nonzero weights, which is equal to the number of columns of the matrix \mathbf{S} .

r such that $k < r \leq n$, there exists a set of weights $s_i \geq 0$ ($i = 1 \dots n$), at most r of which are nonzero, such that

$$\lambda_k \left(\sum_{i=1}^n s_i \mathbf{v}_i \mathbf{v}_i^T \right) \geq \left(1 - \sqrt{\frac{k}{r}} \right)^2, \quad \text{Tr} \left(\sum_{i=1}^n s_i \mathbf{a}_i \mathbf{a}_i^T \right) \leq \text{Tr} \left(\sum_{i=1}^n \mathbf{a}_i \mathbf{a}_i^T \right) = \sum_{i=1}^n \|\mathbf{a}_i\|_2^2.$$

The weights s_i can be computed deterministically in $O(nrk^2 + n\ell)$ time.

In matrix notation (here \mathbf{A} denotes the matrix whose rows are the vectors \mathbf{a}_i), the above lemma guarantees that $\sigma_k(\mathbf{V}^T \mathbf{S}) \geq 1 - \sqrt{k/r}$ and $\|\mathbf{A}^T \mathbf{S}\|_F^2 \leq \|\mathbf{A}\|_F^2$. Observe that the second condition is actually just a general statement about preserving the sum of positive numbers, though for our specific application context, as stated in the lemma, this sum happens to be the Frobenius norm of \mathbf{A} .

4. Proofs of our main results. In this section, we leverage the main tools described in section 3 in order to prove the results of section 1.2 (Theorems 1 through 5). We start with a proof of Theorem 1, using Lemmas 10 and 13.

4.1. Proof of Theorem 1. Apply the algorithm of Lemma 13 on the following two sets of vectors: the n rows of the matrix \mathbf{V}_k and the n rows of the matrix $\mathbf{V}_{\rho-k}$. The output of the algorithm is a sampling and rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times r}$ (see the discussion after Lemma 13 in section 3.2). Let $\mathbf{C} = \mathbf{A}\mathbf{S}$ and note that \mathbf{C} consists of a subset of r rescaled columns of \mathbf{A} . Lemma 13 guarantees that $\sigma_k(\mathbf{V}_k^T \mathbf{S}) \geq 1 - \sqrt{k/r} > 0$ (assuming $r > k$), and so $\text{rank}(\mathbf{V}_k^T \mathbf{S}) = k$. Also, $\sigma_1(\mathbf{V}_{\rho-k}^T \mathbf{S}) = \|\mathbf{V}_{\rho-k}^T \mathbf{S}\|_2 \leq 1 + \sqrt{(\rho-k)/r}$. Applying (4) of Lemma 10, we obtain

$$\begin{aligned} \|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2^2 &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 + \|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}(\mathbf{V}_k^T \mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 + \|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}\|_2^2 \|(\mathbf{V}_k^T \mathbf{S})^+\|_2^2 \\ &= \|\mathbf{A} - \mathbf{A}_k\|_2^2 + \|\mathbf{U}_{\rho-k} \Sigma_{\rho-k} \mathbf{V}_{\rho-k}^T \mathbf{S}\|_2^2 \|(\mathbf{V}_k^T \mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 + \|\Sigma_{\rho-k}\|_2^2 \|\mathbf{V}_{\rho-k}^T \mathbf{S}\|_2^2 \|(\mathbf{V}_k^T \mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 \left(1 + \frac{(1 + \sqrt{(\rho-k)/r})^2}{(1 - \sqrt{k/r})^2} \right), \end{aligned}$$

where the last inequality follows because $\|\Sigma_{\rho-k}\|_2 = \|\mathbf{A} - \mathbf{A}_k\|_2$ and $\|(\mathbf{V}_k^T \mathbf{S})^+\|_2 = 1/\sigma_k(\mathbf{V}_k^T \mathbf{S}) \leq 1/(1 - \sqrt{k/r})$. Theorem 1 now follows by taking square roots of both sides and using $\sqrt{1+x^2} \leq 1+x$. The running time is equal to the time needed to compute \mathbf{V}_k and $\mathbf{V}_{\rho-k}$ plus the running time of the algorithm in Lemma 13. Finally, we note that rescaling the columns of \mathbf{C} does not change the span of its columns and thus is irrelevant in the construction of $\Pi_{\mathbf{C},k}^2(\mathbf{A})$.

Our next theorem describes a deterministic algorithm for spectral norm reconstruction that only needs to compute \mathbf{V}_k and will serve as a prequel to the proof of Theorem 3.

THEOREM 15. *Given $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ and a target rank $k < \rho$, there exists a deterministic polynomial-time algorithm to select $r > k$ columns of \mathbf{A} and form a matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^2(\mathbf{A})\|_2 \leq \frac{1 + \sqrt{n/r}}{1 - \sqrt{k/r}} \cdot \|\mathbf{A} - \mathbf{A}_k\|_2.$$

The matrix \mathbf{C} can be computed in $T_{\mathbf{V}_k} + O(nrk^2)$ time, where $T_{\mathbf{V}_k}$ is the time needed to compute the top k right singular vectors of \mathbf{A} .

Proof. First, apply the algorithm of Lemma 13 on the following two sets of vectors: the n rows of the matrix \mathbf{V}_k and the n rows of the matrix \mathbf{I}_n . The output of the algorithm is a sampling and rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times r}$ (see the discussion after Lemma 13 in section 3.2). Let $\mathbf{C} = \mathbf{A}\mathbf{S}$ and note that \mathbf{C} consists of a subset of r rescaled columns of \mathbf{A} . Lemma 13 guarantees that $\|\mathbf{I}_n\mathbf{S}\|_2 \leq 1 + \sqrt{n/r}$. Applying (5) of Lemma 10, we get

$$\begin{aligned} \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_2^2 &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 \cdot \|\mathbf{S}(\mathbf{V}_k^T\mathbf{S})^+\|_2^2 \\ &= \|\mathbf{A} - \mathbf{A}_k\|_2^2 \cdot \|\mathbf{I}_n\mathbf{S}(\mathbf{V}_k^T\mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 \cdot \|\mathbf{I}_n\mathbf{S}\|_2^2 \cdot \|(\mathbf{V}_k^T\mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{A} - \mathbf{A}_k\|_2^2 \cdot \left(1 + \sqrt{n/r}\right)^2 \cdot \left(1 - \sqrt{k/r}\right)^{-2}. \end{aligned}$$

Again, as in Theorem 1, the rescaling of the columns of \mathbf{C} is irrelevant to the construction of $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$. To analyze the running time of the proposed algorithm, we need to look more closely at Lemma 13 and the related algorithm. The proof of this lemma in section 7 argues that the algorithm of Lemma 13 can be implemented in $O(nrk^2)$ time. The total running time is the time needed to compute \mathbf{V}_k plus $O(nrk^2)$. \square

4.2. Proof of Theorem 3. In order to prove Theorem 3 we will follow the proof of Theorem 15 using Lemma 11 (a fast matrix factorization) instead of Lemma 10 (the exact SVD of \mathbf{A}). More specifically, instead of using the top k right singular vectors of \mathbf{A} (the matrix \mathbf{V}_k), we use the matrix $\mathbf{Z} \in \mathbb{R}^{n \times k}$ of Lemma 11. We now apply the algorithm of Lemma 13 to the following two sets of vectors: the n rows of the matrix \mathbf{Z} and the n rows of the matrix \mathbf{I}_n . The output of the algorithm is a sampling and rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times r}$ (see the discussion after Lemma 13 in section 3.2). Let $\mathbf{C} = \mathbf{A}\mathbf{S}$ and note that \mathbf{C} consists of a subset of r rescaled columns of \mathbf{A} . Applying (2) of Lemma 9, we get

$$\begin{aligned} \|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_2^2 &\leq \|\mathbf{E}\|_2^2 \cdot \|\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\|_2^2 \\ &= \|\mathbf{E}\|_2^2 \cdot \|\mathbf{I}_n\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{E}\|_2^2 \cdot \|\mathbf{I}_n\mathbf{S}\|_2^2 \cdot \|(\mathbf{Z}^T\mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{E}\|_2^2 \cdot \left(1 + \sqrt{n/r}\right)^2 \cdot \left(1 - \sqrt{k/r}\right)^{-2}, \end{aligned}$$

where \mathbf{E} is the residual error from the matrix factorization of Lemma 11. Taking square roots and using the bounds guaranteed by Lemma 13 for $\|\mathbf{I}_n\mathbf{S}\|_2$ and $\|(\mathbf{Z}^T\mathbf{S})^+\|_2$, we obtain a bound in terms of $\|\mathbf{E}\|_2$,

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^\xi(\mathbf{A})\|_2 \leq \|\mathbf{E}\|_2 \cdot \left(1 + \sqrt{n/r}\right) \cdot \left(1 - \sqrt{k/r}\right)^{-1}.$$

Finally, since \mathbf{E} is a random variable, taking expectations and applying the bound of Lemma 11 concludes the proof of the theorem. Again, the rescaling of the columns of \mathbf{C} is irrelevant to the construction of $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$. The running time is the time needed to compute the matrix \mathbf{Z} from Lemma 11 plus an additional $O(nrk^2)$ time as in Theorem 15.

4.3. Proof of Theorem 2. First, apply the algorithm of Lemma 14 to the following two sets of vectors: the n rows of the matrix \mathbf{V}_k and the n rows of the matrix $(\mathbf{A} - \mathbf{A}_k)^T$. The output of the algorithm is a sampling and rescaling matrix

$\mathbf{S} \in \mathbb{R}^{n \times r}$ (see the discussion after Lemma 13 in section 3.2). Let $\mathbf{C} = \mathbf{A}\mathbf{S}$ and note that \mathbf{C} consists of a subset of r rescaled columns of \mathbf{A} . We follow the proof of Theorem 1 in the previous section up to the point where we need to bound the term $\|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}(\mathbf{V}_k^T\mathbf{S})^+\|_{\mathbb{F}}^2$. By strong submultiplicativity,

$$\|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}(\mathbf{V}_k^T\mathbf{S})^+\|_{\mathbb{F}}^2 \leq \|(\mathbf{A} - \mathbf{A}_k)\mathbf{S}\|_{\mathbb{F}}^2 \|(\mathbf{V}_k^T\mathbf{S})^+\|_2^2.$$

To conclude, we apply Lemma 14 to bound the two terms on the right-hand side of the above inequality. The rescaling of the columns of \mathbf{C} is irrelevant to the construction of $\Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})$. The running time of the proposed algorithm is equal to the time needed to compute \mathbf{V}_k plus the time needed to compute $\mathbf{A} - \mathbf{A}_k$ (which is equal to $O(mnk)$ given \mathbf{V}_k) plus the time needed to run the algorithm of Lemma 14, which is equal to $O(nrk^2 + nm)$.

4.4. Proof of Theorem 4. We will follow the proof of Theorem 2, but, as with the proof of Theorem 3, instead of using the top k left singular vectors of \mathbf{A} (the matrix \mathbf{V}_k), we will use the matrix \mathbf{Z} of Lemma 12 that is computed via a fast, approximate matrix factorization. More specifically, let \mathbf{Z} be the matrix of Lemma 12 and run the algorithm of Lemma 14 on the following two sets of vectors: the n rows of the matrix \mathbf{Z} and the n rows of the matrix \mathbf{E}^T . The output of the algorithm is a sampling and rescaling matrix $\mathbf{S} \in \mathbb{R}^{n \times r}$ (see the discussion after Lemma 13 in section 3.2). Let $\mathbf{C} = \mathbf{A}\mathbf{S}$ and note that \mathbf{C} consists of a subset of r rescaled columns of \mathbf{A} . The proof of Theorem 4 is now identical to the proof of Theorem 2, except for using (2) of Lemma 9 instead of (5) in Lemma 10. Ultimately, we obtain

$$\begin{aligned} \|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2 &\leq \|\mathbf{E}\|_{\mathbb{F}}^2 + \|\mathbf{E}\mathbf{S}(\mathbf{Z}^T\mathbf{S})^+\|_2^2 \\ &\leq \|\mathbf{E}\|_{\mathbb{F}}^2 + \|\mathbf{E}\mathbf{S}\|_{\mathbb{F}}^2 \|(\mathbf{Z}^T\mathbf{S})^+\|_2^2 \\ &\leq \left(1 + \left(1 - \sqrt{k/r}\right)^{-2}\right) \|\mathbf{E}\|_{\mathbb{F}}^2. \end{aligned}$$

The last inequality follows from the bounds of Lemma 14. The theorem now follows by taking the expectation of both sides and using Lemma 12 to bound $\mathbb{E}[\|\mathbf{E}\|_{\mathbb{F}}^2]$. The rescaling of the columns of \mathbf{C} is irrelevant to the construction of $\Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})$. The overall running time is derived by replacing the time needed to compute \mathbf{V}_k in Theorem 2 with the time needed to compute the fast approximate factorization of Lemma 12.

4.5. Proof of Theorem 5. Finally, we will prove Theorem 5 by combining the results of Theorem 4 (a constant factor approximation algorithm) with one round of adaptive sampling. We first recall the following lemma, which has appeared in prior work [10, 15].

LEMMA 16. *Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, a target rank k , and an integer r , there exists an algorithm to select r columns from \mathbf{A} to form the matrix $\mathbf{C} \in \mathbb{R}^{m \times r}$ such that*

$$\mathbb{E}[\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2] \leq \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2 + \frac{k}{r} \|\mathbf{A}\|_{\mathbb{F}}^2.$$

The matrix \mathbf{C} can be computed in $O(mn + r \log r)$ time.

Algorithms for the above lemma choose r columns of \mathbf{A} in r independent and identically distributed (i.i.d.) trials, where in each trial a column of \mathbf{A} is sampled with probability proportional to its norm squared (importance sampling). We now state Theorem 2.1 of [7], which builds upon Lemma 16 to provide an adaptive sampling procedure that improves the accuracy guarantees of Lemma 16.

LEMMA 17. Given a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, let $\mathbf{C}_1 \in \mathbb{R}^{m \times r}$ consist of r columns of \mathbf{A} , and define the residual $\mathbf{B} = \mathbf{A} - \mathbf{C}_1 \mathbf{C}_1^+ \mathbf{A} \in \mathbb{R}^{m \times n}$. For $i = 1, \dots, n$, let

$$p_i = \|\mathbf{b}_i\|_2^2 / \|\mathbf{B}\|_F^2,$$

where \mathbf{b}_i is the i th column of the matrix \mathbf{B} . Sample a further s columns from \mathbf{A} in s i.i.d. trials, where in each trial the i th column is chosen with probability p_i . Let $\mathbf{C}_2 \in \mathbb{R}^{m \times s}$ contain the s sampled columns and let $\mathbf{C} = [\mathbf{C}_1 \ \mathbf{C}_2] \in \mathbb{R}^{m \times (r+s)}$ contain the columns of both \mathbf{C}_1 and \mathbf{C}_2 , all of which are columns of \mathbf{A} . Then, for any integer $k > 0$,

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2] \leq \|\mathbf{A} - \mathbf{A}_k\|_F^2 + \frac{k}{s} \|\mathbf{B}\|_F^2.$$

Note that Lemma 17 is an extension of Lemma 16, which can be derived by setting \mathbf{C}_1 to be empty in Lemma 17. We are now ready to prove Theorem 5. First, fix $d > 1$ and define

$$c_0 = (1 + \epsilon_0) \left(1 + 1 / \left(1 - \sqrt{k/\hat{r}} \right)^2 \right),$$

where

$$\hat{r} = \lceil dk \rceil.$$

(We will choose d and ϵ_0 later.) Now run the algorithm of Theorem 4 to sample $\hat{r} = \lceil dk \rceil$ columns of \mathbf{A} and form the matrix \mathbf{C}_1 . Then run the adaptive sampling algorithm of Lemma 17 with $\mathbf{B} = \mathbf{A} - \mathbf{C}_1 \mathbf{C}_1^+ \mathbf{A}$ and sample a further

$$s = \lceil c_0 k / \epsilon \rceil$$

columns of \mathbf{A} to form the matrix \mathbf{C}_2 . Let

$$\mathbf{C} = [\mathbf{C}_1 \ \mathbf{C}_2] \in \mathbb{R}^{n \times (\hat{r}+s)}$$

contain all the sampled columns. We will analyze the expectation

$$\mathbb{E} \|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2.$$

Using the bound of Lemma 17, we first compute the expectation with respect to \mathbf{C}_2 conditioned on \mathbf{C}_1 :

$$\mathbb{E}_{\mathbf{C}_2} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 \mid \mathbf{C}_1] \leq \|\mathbf{A} - \mathbf{A}_k\|_F^2 + \frac{k}{s} \|\mathbf{B}\|_F^2.$$

We now compute the expectation with respect to \mathbf{C}_1 (only \mathbf{B} depends on \mathbf{C}_1):

$$(6) \quad \mathbb{E}_{\mathbf{C}_1} [\mathbb{E}_{\mathbf{C}_2} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2 \mid \mathbf{C}_1]] \leq \|\mathbf{A} - \mathbf{A}_k\|_F^2 + \frac{k}{s} \mathbb{E}_{\mathbf{C}_1} [\|\mathbf{A} - \mathbf{C}_1 \mathbf{C}_1^+ \mathbf{A}\|_F^2].$$

By the law of iterated expectation, the left-hand side is equal to $\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^F(\mathbf{A})\|_F^2]$. We now use the accuracy guarantee of Theorem 4 and our definition of c_0 to bound

$$\mathbb{E}_{\mathbf{C}_1} [\|\mathbf{A} - \mathbf{C}_1 \mathbf{C}_1^+ \mathbf{A}\|_F^2] \leq \mathbb{E}_{\mathbf{C}_1} [\|\mathbf{A} - \Pi_{\mathbf{C}_1,k}^F(\mathbf{A})\|_F^2] \leq c_0 \|\mathbf{A} - \mathbf{A}_k\|_F^2.$$

Using the bound in (6), we obtain

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbf{F}}(\mathbf{A})\|_{\mathbf{F}}^2] \leq \|\mathbf{A} - \mathbf{A}_k\|_{\mathbf{F}}^2 (1 + c_0 k/s).$$

Finally, recall that for our choice of s , $s \geq c_0 k/\epsilon$, and so we obtain the relative-error bound. The number of columns needed is

$$r = \hat{r} + s = dk + c_0 k/\epsilon.$$

Set

$$d = (1 + \alpha)^2,$$

where

$$\alpha = \sqrt[3]{(1 + \epsilon_0)/\epsilon}.$$

After some algebra, this yields

$$r = k \left(\alpha^3 + (1 + \alpha)^3 \right) = \frac{2k}{\epsilon} \left(1 + O \left(\epsilon_0 + \epsilon^{1/3} \right) \right)$$

sampled columns. The time needed to compute the matrix \mathbf{C} is the sum of three terms: the running time of Theorem 4 (which is $O(mnk\epsilon_0^{-1} + n\hat{r}k^2)$), plus the time needed to compute $\mathbf{A} - \mathbf{C}_1 \mathbf{C}_1^+ \mathbf{A}$ (which is $O(mn\hat{r})$), plus the time needed to run the algorithm of Lemma 17 (which is $O(mn + s \log s)$). Assume $r < n$ (otherwise the problem is trivial), set

$$\epsilon_0 = \epsilon^{2/3},$$

and use $d = O(\epsilon^{-2/3})$ to get the final asymptotic run time.

4.5.1. Improving the running time. We conclude by noting that the number of columns required for relative-error approximation is approximately $\frac{2k}{\epsilon}$, a two-factor from optimal, since $\frac{k}{\epsilon}$ columns are necessary (see [9] and section 9.1). We get an improved running time equal to

$$O(mnk + nk^3 + n \log \epsilon^{-1})$$

using just a constant factor more columns by setting d and ϵ_0 in the proof to constants (for example, setting $d = 100$, $\epsilon_0 = \frac{62}{181} \approx \frac{1}{3}$ results in sampling $\frac{3k}{\epsilon}(1 + o(1))$ columns).

5. Proof of Lemma 11: Approximate SVD in the spectral norm. Consider the following algorithm, described in Corollary 10.10 of [19]. The algorithm takes as inputs a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , an integer $2 \leq k < \rho$, an integer $q \geq 1$, and an integer $p \geq 2$. Set $r = k + p$ and construct the matrix $\mathbf{Y} \in \mathbb{R}^{m \times r}$ as follows:

1. Generate an $n \times r$ standard Gaussian matrix \mathbf{R} whose entries are i.i.d. $\mathcal{N}(0, 1)$ variables.
2. Return $\mathbf{Y} = (\mathbf{A}\mathbf{A}^T)^q \mathbf{A}\mathbf{R} \in \mathbb{R}^{m \times r}$.

The running time of the above algorithm is $O(mnrq)$. Corollary 10.10 of [19] presents the following bound:

$$\mathbb{E} [\|\mathbf{A} - \mathbf{Y}\mathbf{Y}^+ \mathbf{A}\|_2] \leq \left(1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{k+p}}{p} \sqrt{\min\{m, n\} - k} \right)^{\frac{1}{2q+1}} \|\mathbf{A} - \mathbf{A}_k\|_2,$$

where $e = 2.718 \dots$. To the best of our understanding, the above result is not immediately applicable to the construction of a factorization of the form $\mathbf{A} = \mathbf{B}\mathbf{Z}^T + \mathbf{E}$, because \mathbf{Y} contains $r > k$ columns. Lemma 26 below, which strengthens Corollary 10.10 in [19], argues that the matrix $\Pi_{\mathbf{Y},k}^2(\mathbf{A})$ contains the desired factorization $\mathbf{B}\mathbf{Z}^T$. Recall that while we cannot compute $\Pi_{\mathbf{Y},k}^2(\mathbf{A})$ efficiently, we can compute a constant-factor approximation, which is sufficient for our purposes. The proof of Lemma 26 is very similar to the proof of Corollary 10.10 of [19], with the only difference being our starting point: instead of using Theorem 9.1 of [19] we use Lemma 10 of our work. To prove Lemma 26, we will need several results for standard Gaussian matrices, projection matrices, and Hölder’s inequality. The following seven lemmas are all borrowed from [19].

LEMMA 18 (Proposition 10.1 in [19]). *Fix matrices \mathbf{X} , \mathbf{Y} , and draw a standard Gaussian matrix \mathbf{R} of appropriate dimensions. Then*

$$\mathbb{E} [\|\mathbf{X}\mathbf{R}\mathbf{Y}\|_2] \leq \|\mathbf{X}\|_2 \|\mathbf{Y}\|_F + \|\mathbf{X}\|_F \|\mathbf{Y}\|_2.$$

LEMMA 19 (Proposition 10.2 in [19]). *For $k, p \geq 2$, draw a standard Gaussian matrix $\mathbf{R} \in \mathbb{R}^{k \times (k+p)}$. Then*

$$\mathbb{E} [\|\mathbf{R}^+\|_2] \leq \frac{e\sqrt{k+p}}{p},$$

where $e = 2.718 \dots$.

LEMMA 20 (Proposition 10.1 in [19]). *Fix matrices \mathbf{X} , \mathbf{Y} , and a standard Gaussian matrix \mathbf{R} of appropriate dimensions. Then*

$$\mathbb{E} [\|\mathbf{X}\mathbf{R}\mathbf{Y}\|_F^2] = \|\mathbf{X}\|_F^2 \|\mathbf{Y}\|_F^2.$$

LEMMA 21 (Proposition 10.2 in [19]). *For $k, p \geq 2$, draw a standard Gaussian matrix $\mathbf{R} \in \mathbb{R}^{k \times (k+p)}$. Then*

$$\mathbb{E} [\|\mathbf{R}^+\|_F^2] = \frac{k}{p-1}.$$

LEMMA 22 (proved in [19]). *For integers $k, p \geq 1$ and a standard Gaussian $\mathbf{R} \in \mathbb{R}^{k \times (k+p)}$ the rank of \mathbf{R} is equal to k with probability one.*

LEMMA 23 (Proposition 8.6 in [19]). *Let \mathbf{P} be a projection matrix. For any matrix \mathbf{X} of appropriate dimensions and an integer $q \geq 0$,*

$$\|\mathbf{P}\mathbf{X}\|_2 \leq \left(\|\mathbf{P}(\mathbf{X}\mathbf{X}^T)^q \mathbf{X}\|_2 \right)^{\frac{1}{2q+1}}.$$

LEMMA 24 (Hölder’s inequality). *Let x be a positive random variable. Then, for any $h \geq 1$,*

$$\mathbb{E} [x] \leq \left(\mathbb{E} [x^h] \right)^{\frac{1}{h}}.$$

The following lemma provides an alternative definition for $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ which will be useful in subsequent proofs. Recall from section 2.2 that we can write $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) = \mathbf{C}\mathbf{X}^\xi$, where

$$\mathbf{X}^\xi = \underset{\Psi \in \mathbb{R}^{r \times n} : \text{rank}(\Psi) \leq k}{\text{argmin}} \|\mathbf{A} - \mathbf{C}\Psi\|_\xi^2.$$

The next lemma basically says that $\Pi_{\mathbf{C},k}^\xi(\mathbf{A})$ is the projection of \mathbf{A} onto the rank- k subspace spanned by $\mathbf{C}\mathbf{X}^\xi$, and that no other subspace in the column space of \mathbf{C} is better.

LEMMA 25. For $\mathbf{A} \in \mathbb{R}^{m \times n}$ and $\mathbf{C} \in \mathbb{R}^{m \times r}$, integer $r > k$, let $\Pi_{\mathbf{C},k}^\xi(\mathbf{A}) = \mathbf{C}\mathbf{X}^\xi$, and let $\mathbf{Y} \in \mathbb{R}^{r \times n}$ be any matrix of rank at most k . Then

$$\|\mathbf{A} - \mathbf{C}\mathbf{X}^\xi\|_\xi^2 = \|\mathbf{A} - (\mathbf{C}\mathbf{X}^\xi)(\mathbf{C}\mathbf{X}^\xi)^+ \mathbf{A}\|_\xi^2 \leq \|\mathbf{A} - (\mathbf{C}\mathbf{Y})(\mathbf{C}\mathbf{Y})^+ \mathbf{A}\|_\xi^2,$$

where $\mathbf{Y} \in \mathbb{R}^{r \times n}$ is any matrix of rank at most k .

Proof. The second inequality will follow from the optimality of \mathbf{X}^ξ because $\mathbf{Y}(\mathbf{C}\mathbf{Y})^+ \mathbf{A}$ has rank at most k . So we only need to prove the first equality. Again, by the optimality of \mathbf{X}^ξ and because $\mathbf{X}^\xi(\mathbf{C}\mathbf{X}^\xi)^+ \mathbf{A}$ has rank at most k ,

$$\|\mathbf{A} - \mathbf{C}\mathbf{X}^\xi\|_\xi^2 \leq \|\mathbf{A} - (\mathbf{C}\mathbf{X}^\xi)(\mathbf{C}\mathbf{X}^\xi)^+ \mathbf{A}\|_\xi^2.$$

To get the reverse inequality, we will use matrix-Pythagoras as follows:

$$\begin{aligned} \|\mathbf{A} - \mathbf{C}\mathbf{X}^\xi\|_\xi^2 &= \left\| \left(\mathbf{I}_m - (\mathbf{C}\mathbf{X}^\xi)(\mathbf{C}\mathbf{X}^\xi)^+ \right) \mathbf{A} - \mathbf{C}\mathbf{X}^\xi \left(\mathbf{I}_n - (\mathbf{C}\mathbf{X}^\xi)^+ \mathbf{A} \right) \right\|_\xi^2 \\ &\geq \left\| \left(\mathbf{I}_m - (\mathbf{C}\mathbf{X}^\xi)(\mathbf{C}\mathbf{X}^\xi)^+ \right) \mathbf{A} \right\|_\xi^2. \quad \square \end{aligned}$$

LEMMA 26 (extension of Corollary 10.10 of [19]). Let \mathbf{A} be a matrix in $\mathbb{R}^{m \times n}$ of rank ρ , let k be an integer satisfying $2 \leq k < \rho$, and let $r = k + p$ for some integer $p \geq 2$. Let $\mathbf{R} \in \mathbb{R}^{n \times r}$ be a standard Gaussian matrix (i.e., a matrix whose entries are drawn in i.i.d. trials from $\mathcal{N}(0, 1)$). Define $\mathbf{B} = (\mathbf{A}\mathbf{A}^\top)^q \mathbf{A}$ and compute $\mathbf{Y} = \mathbf{B}\mathbf{R}$. Then, for any $q \geq 0$,

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2] \leq \left(1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{k+p}}{p} \sqrt{\min\{m,n\} - k} \right)^{\frac{1}{2q+1}} \|\mathbf{A} - \mathbf{A}_k\|_2.$$

Proof. Let $\Pi_{\mathbf{Y},k}^2(\mathbf{A}) = \mathbf{Y}\mathbf{X}_1$ and $\Pi_{\mathbf{Y},k}^2(\mathbf{B}) = \mathbf{Y}\mathbf{X}_2$, where \mathbf{X}_1 is optimal for \mathbf{A} and \mathbf{X}_2 for \mathbf{B} . From Lemma 25,

$$\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2 = \|(\mathbf{I}_m - (\mathbf{Y}\mathbf{X}_1)(\mathbf{Y}\mathbf{X}_1)^+) \mathbf{A}\|_2 \leq \|(\mathbf{I}_m - (\mathbf{Y}\mathbf{X}_2)(\mathbf{Y}\mathbf{X}_2)^+) \mathbf{A}\|_2.$$

From Lemma 23 and using the fact that $\mathbf{I}_m - (\mathbf{Y}\mathbf{X}_2)(\mathbf{Y}\mathbf{X}_2)^+$ is a projection,

$$\begin{aligned} \|(\mathbf{I}_m - (\mathbf{Y}\mathbf{X}_2)(\mathbf{Y}\mathbf{X}_2)^+) \mathbf{A}\|_2 &\leq \left\| \left(\mathbf{I}_m - (\mathbf{Y}\mathbf{X}_2)(\mathbf{Y}\mathbf{X}_2)^+ \right) \left(\mathbf{A}\mathbf{A}^\top \right)^q \mathbf{A} \right\|_2^{\frac{1}{2q+1}} \\ &= \|\mathbf{B} - (\mathbf{Y}\mathbf{X}_2)(\mathbf{Y}\mathbf{X}_2)^+ \mathbf{B}\|_2^{\frac{1}{2q+1}} \\ &= \|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2^{\frac{1}{2q+1}}, \end{aligned}$$

where the last step follows from Lemma 25. We conclude that

$$\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2 \leq \|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2^{\frac{1}{2q+1}}.$$

The matrix \mathbf{Y} is generated using a random \mathbf{R} , so taking expectations and applying Hölder's inequality, we get

$$(7) \quad \mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2] \leq (\mathbb{E} [\|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2])^{\frac{1}{2q+1}}.$$

We now focus on bounding the term on the right-hand side of the above equation. Let the SVD of \mathbf{B} be $\mathbf{B} = \mathbf{U}_\mathbf{B} \boldsymbol{\Sigma}_\mathbf{B} \mathbf{V}_\mathbf{B}^\top$, with the top rank- k factors from the SVD of \mathbf{B} being $\mathbf{U}_{\mathbf{B},k}$, $\boldsymbol{\Sigma}_{\mathbf{B},k}$, and $\mathbf{V}_{\mathbf{B},k}$ and the corresponding trailing factors being $\mathbf{U}_{\mathbf{B},\tau}$, $\boldsymbol{\Sigma}_{\mathbf{B},\tau}$, and $\mathbf{V}_{\mathbf{B},\tau}$. Let $\rho_\mathbf{B}$ be the rank of \mathbf{B} . Let

$$\boldsymbol{\Omega}_1 = \mathbf{V}_{\mathbf{B},k}^\top \mathbf{R} \in \mathbb{R}^{k \times r} \quad \text{and} \quad \boldsymbol{\Omega}_2 = \mathbf{V}_{\mathbf{B},\tau}^\top \mathbf{R} \in \mathbb{R}^{(\rho_\mathbf{B}-k) \times r}.$$

The Gaussian distribution is rotationally invariant, so $\boldsymbol{\Omega}_1, \boldsymbol{\Omega}_2$ are also standard Gaussian matrices which are stochastically independent because $\mathbf{V}_\mathbf{B}^\top$ can be extended to a full rotation. Thus, $\mathbf{V}_{\mathbf{B},k}^\top \mathbf{R}$ and $\mathbf{V}_{\mathbf{B},\tau}^\top \mathbf{R}$ also have entries that are i.i.d. $\mathcal{N}(0,1)$ variables. We now apply Lemma 10 to reconstruct \mathbf{B} , with $\xi = 2$ and $\mathbf{S} = \mathbf{R}$. The rank requirement in Lemma 10 is satisfied because, from Lemma 22, the rank of $\boldsymbol{\Omega}_1$ is equal to k (as it is a standard normal matrix), and thus the matrix \mathbf{R} satisfies the rank assumptions of Lemma 10. We get that

$$\|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2^2 \leq \|\mathbf{B} - \mathbf{B}_k\|_2^2 + \|(\mathbf{B} - \mathbf{B}_k) \mathbf{R} (\mathbf{V}_{\mathbf{B},k}^\top \mathbf{R})^+\|_2^2 \leq \|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_2^2 + \|\boldsymbol{\Sigma}_{\mathbf{B},\tau} \boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^+\|_2^2.$$

Using $\sqrt{x^2 + y^2} \leq x + y$, we conclude that

$$\|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2 \leq \|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_2 + \|\boldsymbol{\Sigma}_{\mathbf{B},\tau} \boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^+\|_2.$$

We now need to take the expectation with respect to $\boldsymbol{\Omega}_1, \boldsymbol{\Omega}_2$. We first take the expectation with respect to $\boldsymbol{\Omega}_2$, conditioning on $\boldsymbol{\Omega}_1$. We then take the expectation with respect to $\boldsymbol{\Omega}_1$. Since only the second term is stochastic, using Lemma 18, we have

$$\mathbf{E}_{\boldsymbol{\Omega}_2} [\|\boldsymbol{\Sigma}_{\mathbf{B},\tau} \boldsymbol{\Omega}_2 \boldsymbol{\Omega}_1^+\|_2 | \boldsymbol{\Omega}_1] \leq \|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_2 \|\boldsymbol{\Omega}_1^+\|_F + \|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_F \|\boldsymbol{\Omega}_1^+\|_2.$$

We now take the expectation with respect to $\boldsymbol{\Omega}_1$. To bound $\mathbb{E} [\|\boldsymbol{\Omega}_1^+\|_2]$, we use Lemma 19. To bound $\mathbb{E} [\|\boldsymbol{\Omega}_1^+\|_F]$, we first use Hölder’s inequality to bound

$$\mathbb{E} [\|\boldsymbol{\Omega}_1^+\|_F] \leq \mathbb{E} [\|\boldsymbol{\Omega}_1^+\|_F^2]^{1/2},$$

and then we use Lemma 21. Since

$$\|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_F \leq \sqrt{\min(m,n) - k} \|\boldsymbol{\Sigma}_{\mathbf{B},\tau}\|_2,$$

collecting our results together, we obtain

$$\mathbb{E} [\|\mathbf{B} - \Pi_{\mathbf{Y},k}^2(\mathbf{B})\|_2] \leq \left(1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{k+p}}{p} \sqrt{\min(m,n) - k} \right) \|\mathbf{B} - \mathbf{B}_k\|_2.$$

To conclude, combine with (7) and note that

$$\|\mathbf{B} - \mathbf{B}_k\|_2 = \|\mathbf{A} - \mathbf{A}_k\|_2^{2q+1}. \quad \square$$

We now have all the necessary ingredients to prove Lemma 11. Let \mathbf{Y} be the matrix of Lemma 26. Set $p = k$ and

$$q = \left\lceil \frac{\log \left(1 + \sqrt{\frac{k}{k-1}} + \frac{e\sqrt{2k}}{k} \sqrt{\min\{m,n\} - k} \right)}{2 \log(1 + \epsilon/\sqrt{2})} \right\rceil - \frac{1}{2},$$

so that

$$\left(1 + \sqrt{\frac{k}{p-1}} + \frac{e\sqrt{k+p}}{p} \sqrt{\min\{m, n\} - k}\right)^{\frac{1}{2q+1}} \leq 1 + \frac{\epsilon}{\sqrt{2}}.$$

Then

$$(8) \quad \mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2] \leq \left(1 + \frac{\epsilon}{\sqrt{2}}\right) \|\mathbf{A} - \mathbf{A}_k\|_2.$$

Given \mathbf{Y} , let \mathbf{Q} be an orthonormal basis for its column space. Then, using the algorithm of section 2.2 and applying Lemma 8, we can construct the matrix $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ such that

$$\|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_2 \leq \sqrt{2} \|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2.$$

Clearly, $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ is a rank- k matrix; let $\mathbf{Z} \in \mathbb{R}^{n \times k}$ denote the matrix containing the right singular vectors of $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$, so $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k = \mathbf{XZ}^T$. Note that \mathbf{Z} is equal to the right singular vectors of the matrix $(\mathbf{Q}^T \mathbf{A})_k$ (because \mathbf{Q} has orthonormal columns), and so \mathbf{Z} has already been computed at the second step of the algorithm of section 2.2. Since $\mathbf{E} = \mathbf{A} - \mathbf{AZZ}^T$ and $\|\mathbf{A} - \mathbf{AZZ}^T\|_2 \leq \|\mathbf{A} - \mathbf{XZ}^T\|_2$ for any \mathbf{X} , we have

$$\|\mathbf{E}\|_2 \leq \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_2 \leq \sqrt{2} \|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_2.$$

Note that, by construction, $\mathbf{EZ} = \mathbf{0}_{m \times k}$. The running time follows by adding the running time of the algorithm at the beginning of this section and the running time of the algorithm of Lemma 2.2.

6. Proof of Lemma 12: Approximate SVD in the Frobenius norm.

Consider the following algorithm, described in Theorem 10.5 of [19]. The algorithm takes as inputs a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ of rank ρ , an integer $2 \leq k < \rho$, and an integer $p \geq 2$. Set $r = k + p$ and construct the matrix $\mathbf{Y} \in \mathbb{R}^{m \times r}$ as follows:

1. Generate an $n \times r$ standard Gaussian matrix \mathbf{R} whose entries are i.i.d. $\mathcal{N}(0, 1)$ variables.
2. Return $\mathbf{Y} = \mathbf{AR} \in \mathbb{R}^{m \times r}$.

The running time of the above algorithm is $O(mnr)$. Theorem 10.5 in [19] presents the following bound:

$$\mathbb{E} [\|\mathbf{A} - \mathbf{YY}^+ \mathbf{A}\|_{\mathbb{F}}] \leq \left(1 + \frac{k}{p-1}\right)^{\frac{1}{2}} \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}.$$

To the best of our understanding, the above result is not immediately applicable to the construction of a factorization of the form $\mathbf{A} = \mathbf{BZ}^T + \mathbf{E}$ (as in Lemma 12) because \mathbf{Y} contains $r > k$ columns.

LEMMA 27 (extension of Theorem 10.5 of [19]). *Let \mathbf{A} be a matrix in $\mathbb{R}^{m \times n}$ of rank ρ , let k be an integer satisfying $2 \leq k < \rho$, and let $r = k + p$ for some integer $p \geq 2$. Let $\mathbf{R} \in \mathbb{R}^{n \times r}$ be a standard Gaussian matrix (i.e., a matrix whose entries are drawn in i.i.d. trials from $\mathcal{N}(0, 1)$) and compute $\mathbf{Y} = \mathbf{AR}$. Then*

$$\mathbb{E} [\|\mathbf{A} - \Pi_{\mathbf{Y},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2] \leq \left(1 + \frac{k}{p-1}\right) \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2.$$

Proof. We construct the matrix \mathbf{Y} as described at the beginning of this section. Let the rank of \mathbf{A} be ρ and let $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ be the SVD of \mathbf{A} . Define

$$\mathbf{\Omega}_1 = \mathbf{V}_k^T \mathbf{R} \in \mathbb{R}^{k \times r} \quad \text{and} \quad \mathbf{\Omega}_2 = \mathbf{V}_{\rho-k}^T \mathbf{R} \in \mathbb{R}^{(\rho-k) \times r}.$$

The Gaussian distribution is rotationally invariant, so $\mathbf{\Omega}_1, \mathbf{\Omega}_2$ are also standard Gaussian matrices which are stochastically independent because \mathbf{V}^T can be extended to a full rotation. Thus, $\mathbf{V}_k^T \mathbf{R}$ and $\mathbf{V}_{\rho-k}^T \mathbf{R}$ also have entries that are i.i.d. $\mathcal{N}(0, 1)$ variables. We now apply Lemma 10 to reconstructing \mathbf{A} , with $\xi = F$ and $\mathbf{S} = \mathbf{R}$. Recall that from Lemma 22, the rank of $\mathbf{\Omega}_1$ is equal to k , and thus the matrix \mathbf{R} satisfies the rank assumptions of Lemma 10. We have that

$$\|\mathbf{A} - \Pi_{\mathbf{Y},k}^F(\mathbf{A})\|_F^2 \leq \|\mathbf{A} - \mathbf{A}_k\|_F^2 + \|\mathbf{\Sigma}_{\rho-k} \mathbf{\Omega}_2 \mathbf{\Omega}_1^+\|_F^2,$$

where $\mathbf{A} - \mathbf{A}_k = \mathbf{U}_{\rho-k} \mathbf{\Sigma}_{\rho-k} \mathbf{V}_{\rho-k}^T$. To conclude, we take the expectation on both sides, and since only the second term on the right-hand side is stochastic, we bound as follows:

$$\begin{aligned} \mathbb{E} [\|\mathbf{\Sigma}_{\rho-k} \mathbf{\Omega}_2 \mathbf{\Omega}_1^+\|_F^2] &\stackrel{(a)}{=} \mathbf{E}_{\mathbf{\Omega}_1} [\mathbf{E}_{\mathbf{\Omega}_2} [\|\mathbf{\Sigma}_{\rho-k} \mathbf{\Omega}_2 \mathbf{\Omega}_1^+\|_F^2 | \mathbf{\Omega}_1]] \\ &\stackrel{(b)}{=} \mathbf{E}_{\mathbf{\Omega}_1} [\|\mathbf{\Sigma}_{\rho-k}\|_F^2 \|\mathbf{\Omega}_1^+\|_F^2] \\ &\stackrel{(c)}{=} \|\mathbf{\Sigma}_{\rho-k}\|_F^2 \mathbb{E} [\|\mathbf{\Omega}_1^+\|_F^2] \\ &\stackrel{(d)}{=} \frac{k}{p-1} \|\mathbf{\Sigma}_{\rho-k}\|_F^2. \end{aligned}$$

(a) follows from the law of iterated expectation; (b) follows from Lemma 20; (c) follows because $\|\mathbf{\Sigma}_{\rho-k}\|_F^2$ is a constant; (d) follows from Lemma 21. We conclude the proof by noting that $\|\mathbf{\Sigma}_{\rho-k}\|_F = \|\mathbf{A} - \mathbf{A}_k\|_F$. \square

We now have all the necessary ingredients to conclude the proof of Lemma 12. Let \mathbf{Y} be the matrix of Lemma 27, and let \mathbf{Q} be an orthonormal basis for its column space. Then, using the algorithm of section 2.2 and applying Lemma 8, we can construct the matrix $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ such that

$$\|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_F^2 = \|\mathbf{A} - \Pi_{\mathbf{Y},k}^F(\mathbf{A})\|_F^2.$$

Clearly, $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$ is a rank- k matrix; let $\mathbf{Z} \in \mathbb{R}^{n \times k}$ be the matrix containing the right singular vectors of $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k$, so $\mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k = \mathbf{X}\mathbf{Z}^T$. Note that \mathbf{Z} is equal to the right singular vectors of the matrix $(\mathbf{Q}^T \mathbf{A})_k$ (because \mathbf{Q} has orthonormal columns), and thus \mathbf{Z} has already been computed at the second step of the algorithm of section 2.2. Since $\mathbf{E} = \mathbf{A} - \mathbf{A}\mathbf{Z}\mathbf{Z}^T$ and $\|\mathbf{A} - \mathbf{A}\mathbf{Z}\mathbf{Z}^T\|_F \leq \|\mathbf{A} - \mathbf{X}\mathbf{Z}^T\|_F$ for any \mathbf{X} , we have

$$\|\mathbf{E}\|_F^2 \leq \|\mathbf{A} - \mathbf{Q}(\mathbf{Q}^T \mathbf{A})_k\|_F^2 = \|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_F^2.$$

To conclude, take expectations on both sides, use Lemma 27 to bound the term $\mathbb{E}[\|\mathbf{A} - \Pi_{\mathbf{Y},k}^2(\mathbf{A})\|_F^2]$, and set $p = \lceil \frac{k}{\epsilon} + 1 \rceil$ to obtain

$$\mathbb{E} [\|\mathbf{E}\|_F^2] \leq \left(1 + \frac{k}{p-1}\right) \|\mathbf{A} - \mathbf{A}_k\|_F^2 \leq (1 + \epsilon) \|\mathbf{A} - \mathbf{A}_k\|_F^2.$$

By construction, $\mathbf{E}\mathbf{Z} = \mathbf{0}_{m \times k}$. The running time follows by adding the running time of the algorithm at the beginning of this section and the running time of the algorithm of Lemma 2.2.

Algorithm 1 Deterministic dual set spectral sparsification.

Input:

- $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$, with $\sum_{i=1}^n \mathbf{v}_i \mathbf{v}_i^T = \mathbf{I}_k$ ($k < n$)
- $\mathcal{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_n\}$, with $\sum_{i=1}^n \mathbf{u}_i \mathbf{u}_i^T = \mathbf{I}_\ell$ ($\ell \leq n$)
- integer r , with $k < r < n$

Output: A vector of weights $\mathbf{s} = [s_1, \dots, s_n]$, with $s_i \geq 0$ and at most r nonzero s_i 's.

1. Initialize $\mathbf{s}_0 = \mathbf{0}_{n \times 1}$, $\mathbf{A}_0 = \mathbf{0}_{k \times k}$, $\mathbf{B}_0 = \mathbf{0}_{\ell \times \ell}$.
2. For $\tau = 0, \dots, r - 1$

- Compute L_τ and U_τ from (12).
- Find an index j in $\{1, \dots, n\}$ such that

$$(9) \quad U(\mathbf{u}_j, \delta_U, \mathbf{B}_\tau, U_\tau) \leq L(\mathbf{v}_j, \delta_L, \mathbf{A}_\tau, L_\tau).$$

- Let

$$(10) \quad t^{-1} = \frac{1}{2} (U(\mathbf{u}_j, \delta_U, \mathbf{B}_\tau, U_\tau) + L(\mathbf{v}_j, \delta_L, \mathbf{A}_\tau, L_\tau)).$$

- Update the j th component of \mathbf{s} , \mathbf{A}_τ , and \mathbf{B}_τ :

$$(11) \quad \mathbf{s}_{\tau+1}[j] = \mathbf{s}_\tau[j] + t, \quad \mathbf{A}_{\tau+1} = \mathbf{A}_\tau + t \mathbf{v}_j \mathbf{v}_j^T, \quad \text{and} \quad \mathbf{B}_{\tau+1} = \mathbf{B}_\tau + t \mathbf{u}_j \mathbf{u}_j^T.$$

3. Return $\mathbf{s} = r^{-1} (1 - \sqrt{k/r}) \cdot \mathbf{s}_\tau$.
-

7. Dual set spectral sparsification: Proof of Lemma 13. In this section, we prove Lemma 13, which generalizes Theorem 3.1 in [2]. Indeed, setting $\mathcal{V} = \mathcal{U}$ reproduces the spectral sparsification result of Theorem 3.1 in [2]. The basic observation is that the analyses of [2] for the upper bound and lower bound are essentially independent. This means that the analysis in [2] *directly* applies when the upper and lower bounds are analyzed with respect to *different* sets of vectors. The only point at which the bounds are used together is when one needs to compute a weight that is in between the two bounds, which, as [2] shows, is always possible with a single set of vectors, and it is also true with different sets of vectors for the same reason. For completeness we present the details, simplifying a little and emphasizing the independent treatment of \mathbf{v}_i and \mathbf{u}_i , in the proof.

As in [2], we will provide a constructive proof of the lemma and start by describing the algorithm that computes the weights s_i , $i = 1, \dots, n$.

7.1. The algorithm. The fundamental idea underlying Algorithm 1 is the greedy selection of vectors that satisfy a number of desired properties in each step. These properties will eventually imply the eigenvalue bounds of Lemma 13. We start by defining several quantities that will be used in the description of the algorithm and its proof. First, fix two constants:

$$\delta_L = 1; \quad \delta_U = \frac{1 + \sqrt{\frac{\ell}{r}}}{1 - \sqrt{\frac{k}{r}}}.$$

Given k , ℓ , and r (all inputs of Algorithm 1) and a parameter $\tau = 0, \dots, r - 1$, define two parameters L_τ and U_τ as follows:

$$(12) \quad L_\tau = r \left(\frac{\tau}{r} - \sqrt{\frac{k}{r}} \right) = \tau - \sqrt{rk}; \quad U_\tau = \frac{(\tau - r) \left(1 + \sqrt{\frac{\ell}{r}} \right) + r \left(1 + \sqrt{\frac{\ell}{r}} \right)^2}{1 - \sqrt{\frac{k}{r}}} = \delta_U \left(\tau + \sqrt{\ell r} \right).$$

We next define the lower and upper functions $\underline{\phi}(L, \mathbf{A})$ ($L \in \mathbb{R}$ and $\mathbf{A} \in \mathbb{R}^{k \times k}$) and $\overline{\phi}(U, \mathbf{B})$ ($U \in \mathbb{R}$ and $\mathbf{B} \in \mathbb{R}^{\ell \times \ell}$) as follows:

$$(13) \quad \underline{\phi}(L, \mathbf{A}) = \sum_{i=1}^k \frac{1}{\lambda_i(\mathbf{A}) - L}; \quad \overline{\phi}(U, \mathbf{B}) = \sum_{i=1}^{\ell} \frac{1}{U - \lambda_i(\mathbf{B})}.$$

Let $L(\mathbf{v}, \delta_L, \mathbf{A}, L)$ be a function with four inputs (a vector $\mathbf{v} \in \mathbb{R}^{k \times 1}$, $\delta_L \in \mathbb{R}$, a matrix $\mathbf{A} \in \mathbb{R}^{k \times k}$, and $L \in \mathbb{R}$):

$$(14) \quad L(\mathbf{v}, \delta_L, \mathbf{A}, L) = \frac{\mathbf{v}^T(\mathbf{A} - (L + \delta_L)\mathbf{I}_k)^{-2}\mathbf{v}}{\underline{\phi}(L + \delta_L, \mathbf{A}) - \underline{\phi}(L, \mathbf{A})} - \mathbf{v}^T(\mathbf{A} - (L + \delta_L)\mathbf{I}_k)^{-1}\mathbf{v}.$$

Similarly, let $U(\mathbf{u}, \delta_U, \mathbf{B}, U)$ be a function with four inputs (a vector $\mathbf{u} \in \mathbb{R}^{\ell \times 1}$, $\delta_U \in \mathbb{R}$, a matrix $\mathbf{B} \in \mathbb{R}^{\ell \times \ell}$, and $U \in \mathbb{R}$):

$$(15) \quad U(\mathbf{u}, \delta_U, \mathbf{B}, U) = \frac{\mathbf{u}^T((U + \delta_U)\mathbf{I}_{\ell} - \mathbf{B})^{-2}\mathbf{u}}{\overline{\phi}(U, \mathbf{B}) - \overline{\phi}(U + \delta_U, \mathbf{B})} + \mathbf{u}^T((U + \delta_U)\mathbf{I}_{\ell} - \mathbf{B})^{-1}\mathbf{u}.$$

Algorithm 1 runs in r steps. The initial vector of weights \mathbf{s}_0 is initialized to the all-zero vector. At each step $\tau = 0, \dots, r - 1$, the algorithm selects a pair of vectors $(\mathbf{u}_j, \mathbf{v}_j)$ that satisfy (9), computes the associated weight t from (10), and updates two matrices and the vector of weights appropriately, as specified in (11).

7.2. Running time. The algorithm runs in r iterations. In each iteration, we evaluate the functions $U(\mathbf{u}, \delta_U, \mathbf{B}, U)$ and $L(\mathbf{v}, \delta_L, \mathbf{A}, L)$ at most n times. Note that all n evaluations for both functions need at most $O(k^3 + nk^2 + \ell^3 + n\ell^2)$ time, because the matrix inversions can be performed once for all n evaluations. Finally, the updating step needs an additional $O(k^2 + \ell^2)$ time. Overall, the complexity of the algorithm is of the order $O(r(k^3 + nk^2 + \ell^3 + n\ell^2 + k^2 + \ell^2)) = O(rn(k^2 + \ell^2))$.

Note that when \mathcal{U} is the standard basis ($\mathcal{U} = \{\mathbf{e}_1, \dots, \mathbf{e}_n\}$ and $\ell = n$), the computations can be done much more efficiently: the eigenvalues of $\mathbf{B}_{\mathbf{s}}$ need not be computed explicitly (the expensive step), since they are available by inspection, being equal to the weights \mathbf{s}_{τ} . In the function $U(\mathbf{u}, \delta_U, \mathbf{B}, U)$, the functions $\overline{\phi}$ (given the eigenvalues) need only be computed once per iteration, in $O(n)$ time. The remaining terms can be computed in $O(1)$ time, because, for example, $\mathbf{e}_i^T((U + \delta_U)\mathbf{I} - \mathbf{B})^{-2}\mathbf{e}_i = (U + \delta_U - \mathbf{s}[i])^{-2}$. The running time now drops to $O(rnk^2)$, since all the operations on \mathcal{U} only contribute $O(rn)$.

7.3. Proof of correctness. We prove that the output of Algorithm 1 satisfies Lemma 13. Our proof is similar to the proof of Theorem 3.1 of [2]. The main difference is that we need to accommodate two different sets of vectors. Let $\mathbf{W} \in \mathbb{R}^{m \times m}$ be a positive semidefinite matrix with eigendecomposition

$$\mathbf{W} = \sum_{i=1}^m \lambda_i(\mathbf{W})\mathbf{u}_i\mathbf{u}_i^T$$

and recall the functions $\underline{\phi}(L, \mathbf{W})$, $\overline{\phi}(U, \mathbf{W})$, $L(\mathbf{v}, \delta_L, \mathbf{W}, L)$, and $U(\mathbf{v}, \delta_U, \mathbf{W}, U)$ defined in (13), (14), and (15). We now quote two lemmas proven in [2] using the Sherman–Morrison–Woodbury identity. These lemmas allow us to control the smallest and largest eigenvalues of \mathbf{W} under a rank-one perturbation.

LEMMA 28. Fix $\delta_L > 0$, $\mathbf{W} \in \mathbb{R}^{m \times m}$, $\mathbf{v} \in \mathbb{R}^m$, and $L < \lambda_m(\mathbf{W})$. If $t > 0$ satisfies

$$t^{-1} \leq L(\mathbf{v}, \delta_L, \mathbf{W}, L),$$

then $\lambda_m(\mathbf{W} + t\mathbf{v}\mathbf{v}^T) \geq L + \delta_L$.

LEMMA 29. Fix $\delta_U > 0$, $\mathbf{W} \in \mathbb{R}^{m \times m}$, $\mathbf{v} \in \mathbb{R}^m$, and $U > \lambda_1(\mathbf{W})$. If t satisfies

$$t^{-1} \geq U(\mathbf{v}, \delta_U, \mathbf{W}, U),$$

then $\lambda_1(\mathbf{W} + t\mathbf{v}\mathbf{v}^T) \leq U + \delta_U$.

Now recall that Algorithm 1 runs in r steps. Initially, all n weights are set to zero. Assume that at the τ th step ($\tau = 0, \dots, r-1$) the vector of weights

$$\mathbf{s}_\tau = [\mathbf{s}_\tau[1], \dots, \mathbf{s}_\tau[n]]$$

has been constructed, and let

$$\mathbf{A}_\tau = \sum_{i=1}^n \mathbf{s}_\tau[i] \mathbf{v}_i \mathbf{v}_i^T \quad \text{and} \quad \mathbf{B}_\tau = \sum_{i=1}^n \mathbf{s}_\tau[i] \mathbf{u}_i \mathbf{u}_i^T.$$

Note that both matrices \mathbf{A}_τ and \mathbf{B}_τ are positive semidefinite. We claim the following lemma, which guarantees that the algorithm is well-defined. The proof is deferred to the next subsection.

LEMMA 30. At the τ th step, for all $\tau = 0, \dots, r-1$, there exists an index j in $\{1, \dots, n\}$ such that setting the weight $t > 0$ as in (10) satisfies

$$(16) \quad U(\mathbf{u}_j, \delta_U, \mathbf{B}_\tau, U_\tau) \leq t^{-1} \leq L(\mathbf{v}_j, \delta_L, \mathbf{A}_\tau, L_\tau).$$

Once an index j and a weight $t > 0$ have been computed, Algorithm 1 updates the j th weight in the vector of weights \mathbf{s}_τ to create the vector of weights $\mathbf{s}_{\tau+1}$. Clearly, at each of the r steps, only one element of the vector of weights is updated. Since \mathbf{s}_0 is initialized to the all-zeros vector, after all r steps are completed, at most r weights are nonzero. The following lemma argues that $\lambda_{\min}(\mathbf{A}_\tau)$ and $\lambda_{\max}(\mathbf{B}_\tau)$ are bounded.

LEMMA 31. For all $\tau = 0, \dots, r$, $\lambda_{\min}(\mathbf{A}_\tau) \geq L_\tau$ and $\lambda_{\max}(\mathbf{B}_\tau) \leq U_\tau$.

Proof. Recall (12) and observe that $L_0 = -\sqrt{rk} < 0$ and $U_0 = \delta_U \sqrt{r\ell} > 0$. Thus, the lemma holds at $\tau = 0$. It is also easy to verify that $L_{\tau+1} = L_\tau + \delta_L$ and, similarly, $U_{\tau+1} = U_\tau + \delta_U$. Now, at the τ -step, given an index j and a corresponding weight $t > 0$ satisfying (16), Lemmas 28 and 29 imply that

$$\begin{aligned} \lambda_{\min}(\mathbf{A}_{\tau+1}) &= \lambda_{\min}(\mathbf{A}_\tau + t\mathbf{v}_j \mathbf{v}_j^T) \geq L_\tau + \delta_L = L_{\tau+1}; \\ \lambda_{\max}(\mathbf{B}_{\tau+1}) &= \lambda_{\max}(\mathbf{B}_\tau + t\mathbf{u}_j \mathbf{u}_j^T) \leq U_\tau + \delta_U = U_{\tau+1}. \end{aligned}$$

The lemma now follows by simple induction on τ . \square

We are now ready to conclude the proof of Lemma 13. By Lemma 31, at the r th step,

$$\lambda_{\max}(\mathbf{B}_r) \leq U_r \quad \text{and} \quad \lambda_{\min}(\mathbf{A}_r) \geq L_r.$$

Recall the definitions of U_r and L_r from (12) and note that they are both positive and well-defined because $r > k$. Lemma 13 now follows after rescaling the vector of weights \mathbf{s} by $r^{-1}(1 - \sqrt{k/r})$. Note that the rescaling does not change the number of nonzero elements of \mathbf{s} , but does rescale all the eigenvalues of \mathbf{A}_r and \mathbf{B}_r .

7.4. Proof of Lemma 30. In order to prove Lemma 30 we will use the following averaging argument.

LEMMA 32. *At any step $\tau = 0, \dots, r-1$,*

$$\sum_{i=1}^n U(\mathbf{u}_i, \delta_U, \mathbf{B}_\tau, U_\tau) \leq 1 - \sqrt{\frac{k}{r}} \leq \sum_{i=1}^n L(\mathbf{v}_i, \delta_L, \mathbf{A}_\tau, L_\tau).$$

Proof. For notational convenience, let $\bar{\phi}_\tau = \bar{\phi}(U_\tau, \mathbf{B}_\tau)$ and let $\underline{\phi}_\tau = \underline{\phi}(L_\tau, \mathbf{A}_\tau)$. At $\tau = 0$, $\mathbf{B}_0 = \mathbf{0}$ and $\mathbf{A}_0 = \mathbf{0}$ and thus $\bar{\phi}_0 = \ell/U_0$ and $\underline{\phi}_0 = -k/L_0$. Focus on the τ th step and assume that the algorithm has run correctly up to that point. Then $\bar{\phi}_\tau \leq \bar{\phi}_0$ and $\underline{\phi}_\tau \leq \underline{\phi}_0$. Both are true at $\tau = 0$ and, assuming that the algorithm has run correctly until the τ th step, Lemmas 28 and 29 guarantee that $\bar{\phi}_\tau$ and $\underline{\phi}_\tau$ are nonincreasing.

First, consider the upper bound on U . In the following derivation, λ_i denotes the i th eigenvalue of \mathbf{B}_τ . Using $\text{Tr}(\mathbf{u}^T \mathbf{X} \mathbf{u}) = \text{Tr}(\mathbf{X} \mathbf{u} \mathbf{u}^T)$ and $\sum_i \mathbf{u}_i \mathbf{u}_i^T = \mathbf{I}_\ell$, we get

$$\begin{aligned} \sum_{i=1}^n U(\mathbf{u}_i, \delta_U, \mathbf{B}_\tau, U_\tau) &= \frac{\text{Tr}[(U_{\tau+1} \mathbf{I}_\ell - \mathbf{B}_\tau)^{-2}]}{\bar{\phi}_\tau - \bar{\phi}(U_{\tau+1}, \mathbf{B}_\tau)} + \bar{\phi}(U_{\tau+1}, \mathbf{B}_\tau) \\ &= \frac{\sum_{i=1}^\ell \frac{1}{(U_{\tau+1} - \lambda_i)^2}}{\delta_U \sum_{i=1}^\ell \frac{1}{(U_{\tau+1} - \lambda_i)(U_\tau - \lambda_i)}} + \sum_{i=1}^\ell \frac{1}{(U_{\tau+1} - \lambda_i)} \\ &= \frac{1}{\delta_U} + \bar{\phi}_\tau - \frac{1}{\delta_U} \left(1 - \frac{\sum_{i=1}^\ell \frac{1}{(U_{\tau+1} - \lambda_i)^2}}{\sum_{i=1}^\ell \frac{1}{(U_{\tau+1} - \lambda_i)(U_\tau - \lambda_i)}} \right) \\ &\quad - \delta_U \sum_{i=1}^\ell \frac{1}{(U_\tau - \lambda_i)(U_{\tau+1} - \lambda_i)} \\ &\leq \frac{1}{\delta_U} + \bar{\phi}_0. \end{aligned}$$

The last line follows because the last two terms are negative (using the fact that $U_{\tau+1} > U_\tau > \lambda_i$) and $\bar{\phi}_\tau \leq \bar{\phi}_0$. Now, using $\bar{\phi}_0 = \delta_U \sqrt{r\ell}$ and the definition of δ_U , the upper bound follows:

$$\frac{1}{\delta_U} + \bar{\phi}_0 = \frac{1}{\delta_U} + \frac{\ell}{\delta_U \sqrt{r\ell}} = \frac{1}{\delta_U} \left(1 + \sqrt{\frac{\ell}{r}} \right) = 1 - \sqrt{\frac{k}{r}}.$$

In order to prove the lower bound on L we use a similar argument. Let λ_i denote the i th eigenvalue of \mathbf{A}_τ . Then

$$\begin{aligned} \sum_{i=1}^n L(\mathbf{v}_i, \delta_L, \mathbf{A}_\tau, L_\tau) &= \frac{\text{Tr}[(\mathbf{A}_\tau - L_{\tau+1} \mathbf{I}_k)^{-2}]}{\underline{\phi}(L_{\tau+1}, \mathbf{A}_\tau) - \underline{\phi}_\tau} - \underline{\phi}(L_{\tau+1}, \mathbf{A}_\tau) \\ &= \frac{\sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2}}{\delta_L \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})(\lambda_i - L_\tau)}} - \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})} \\ &= \frac{1}{\delta_L} - \underline{\phi}_\tau + \frac{1}{\delta_L} \left(\frac{\sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2}}{\sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})(\lambda_i - L_\tau)}} - 1 \right) \\ &\quad - \delta_L \sum_{i=1}^k \frac{1}{(\lambda_i - L_\tau)(\lambda_i - L_{\tau+1})} \\ &\geq \frac{1}{\delta_L} - \underline{\phi}_0 + \mathcal{E}. \end{aligned}$$

Assuming $\mathcal{E} \geq 0$ the claim follows immediately because $\delta_L = 1$ and

$$\underline{\phi}_0 = \frac{-k}{L_0} = \frac{k}{\sqrt{rk}} = \sqrt{\frac{k}{r}}.$$

Thus, we only need to show that $\mathcal{E} \geq 0$. From the Cauchy–Schwarz inequality, for $a_i, b_i \geq 0$, $(\sum_i a_i b_i)^2 \leq (\sum_i a_i^2 b_i) (\sum_i b_i)$ and thus

$$\begin{aligned} \mathcal{E} \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})(\lambda_i - L_\tau)} &= \frac{1}{\delta_L} \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2(\lambda_i - L_\tau)} \\ (17) \quad &- \delta_L \left(\sum_{i=1}^k \frac{1}{(\lambda_i - L_\tau)(\lambda_i - L_{\tau+1})} \right)^2 \\ &\geq \frac{1}{\delta_L} \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2(\lambda_i - L_\tau)} \\ (18) \quad &- \delta_L \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2(\lambda_i - L_\tau)} \sum_{i=1}^k \frac{1}{\lambda_i - L_\tau} \\ (19) \quad &= \left(\frac{1}{\delta_L} - \delta_L \underline{\phi}_\tau \right) \sum_{i=1}^k \frac{1}{(\lambda_i - L_{\tau+1})^2(\lambda_i - L_\tau)}. \end{aligned}$$

To conclude our proof, first note that

$$\delta_L^{-1} - \delta_L \underline{\phi}_\tau \geq \delta_L^{-1} - \delta_L \underline{\phi}_0 = 1 - \sqrt{\frac{k}{r}} > 0$$

(recall that $r > k$). Second, $\lambda_i > L_{\tau+1}$ because

$$\lambda_{\min}(\mathbf{A}_\tau) > L_\tau + \frac{1}{\underline{\phi}_\tau} \geq L_\tau + \frac{1}{\underline{\phi}_0} = L_\tau + \sqrt{\frac{r}{k}} > L_\tau + 1 = L_{\tau+1}.$$

Combining these two observations with (19) we conclude that $\mathcal{E} \geq 0$. \square

Lemma 30 follows from Lemma 32 because the two inequalities must hold simultaneously for at least one index j .

8. Dual set spectral Frobenius sparsification: Proof of Lemma 14.

In this section we will provide a constructive proof of Lemma 14. Our proof closely follows the proof of Lemma 13, so we will only highlight the differences. We first discuss modifications to Algorithm 1. First of all, the new inputs are $\mathcal{V} = \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ and $\mathcal{A} = \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$. The output is a set of n nonnegative weights s_i , at most r of which are nonzero. We define the parameters

$$\delta_L = 1; \quad \delta_U = \frac{\sum_{i=1}^n \|\mathbf{a}_i\|_2^2}{1 - \sqrt{\frac{k}{r}}}; \quad L_\tau = \tau - \sqrt{rk}; \quad U_\tau = \tau \delta_U$$

for all $\tau = 0, \dots, r - 1$. Let \mathbf{s}_τ denote the vector of weights at the τ th step of Algorithm 1 and initialize \mathbf{s}_0 and \mathbf{A}_0 as in Algorithm 1 (\mathbf{B}_0 will not be necessary). We now define the function $U_F(\mathbf{a}, \delta_U)$, where $\mathbf{a} \in \mathbb{R}^\ell$ and $\delta_U \in \mathbb{R}$:

$$(20) \quad U_F(\mathbf{a}, \delta_U) = \delta_U^{-1} \mathbf{a}^\top \mathbf{a}.$$

Then, at the τ th step, the algorithm will pick an index j and compute a weight $t > 0$ such that

$$(21) \quad U_F(\mathbf{a}_j, \delta_U) \leq t^{-1} \leq L(\mathbf{v}_j, \delta_L, \mathbf{A}_\tau, L_\tau).$$

The algorithm updates the vector of weights \mathbf{s}_τ and the matrix

$$\mathbf{A}_\tau = \sum_{i=1}^n s_{\tau,i} \mathbf{v}_i \mathbf{v}_i^\top.$$

It is worth noting that the algorithm does not need to update the matrix

$$\mathbf{B}_\tau = \sum_{i=1}^n s_{\tau,i} \mathbf{a}_i \mathbf{a}_i^\top,$$

because the function U_F does not need \mathbf{B}_τ as input. To prove the correctness of the algorithm we need the following two intermediate lemmas.

LEMMA 33. *At every step $\tau = 0, \dots, r-1$ there exists an index j in $\{1, \dots, n\}$ that satisfies (21).*

Proof. The proof is very similar to the proof of Lemma 30 (via Lemma 32) so we only sketch the differences. First, note that the dynamics of L have not been changed and thus the lower bound for the average of $L(\mathbf{v}_j, \delta_L, \mathbf{A}_\tau, L_\tau)$ still holds. We only need to upper bound the average of $U_F(\mathbf{a}_i, \delta_U)$ as in Lemma 32. Indeed,

$$\sum_{i=1}^n U_F(\mathbf{a}_i, \delta_U) = \delta_U^{-1} \sum_{i=1}^n \mathbf{a}_i^\top \mathbf{a}_i = \delta_U^{-1} \sum_{i=1}^n \|\mathbf{a}_i\|_2^2 = 1 - \sqrt{\frac{k}{r}},$$

where the last equality follows from the definition of δ_U . \square

LEMMA 34. *Let $W \in \mathbb{R}^{\ell \times \ell}$ be a symmetric positive semidefinite matrix, let $\mathbf{a} \in \mathbb{R}^\ell$ be a vector, and let $U \in \mathbb{R}$ satisfy $U > \text{Tr}(W)$. If $t > 0$ satisfies*

$$U_F(\mathbf{a}, \delta_U) \leq t^{-1},$$

then

$$\text{Tr}(W + t\mathbf{v}\mathbf{v}^\top) \leq U + \delta_U.$$

Proof. Using the conditions of the lemma and the definition of U_F from (20),

$$\begin{aligned} \text{Tr}(W + t\mathbf{a}\mathbf{a}^\top) - U - \delta_U &= \text{Tr}(W) - U + t\mathbf{a}^\top \mathbf{a} - \delta_U, \\ &\leq \text{Tr}(W) - U < 0, \end{aligned}$$

which concludes the proof of the lemma. \square

We can now combine Lemmas 28 and 34 to prove that for all $\tau = 0, \dots, r$,

$$\lambda_{\min}(\mathbf{A}_\tau) \geq L_\tau \quad \text{and} \quad \text{Tr}(\mathbf{B}_\tau) \leq U_\tau.$$

Note that after all r steps of the algorithm are completed,

$$L_r = r \left(1 - \sqrt{k/r}\right)$$

and

$$U_r = r \left(1 - \sqrt{k/r}\right)^{-1} \sum_{i=1}^n \|\mathbf{a}_i\|_2^2.$$

A simple rescaling now concludes the proof. The running time of the (modified) Algorithm 1 is $O(nrk^2 + n\ell)$, where the latter term emerges from the need to compute the function $U_F(\mathbf{a}_j, \delta_U)$ for all $j = 1, \dots, n$ once throughout the algorithm.

9. Lower bounds.

9.1. Spectral norm approximation. Theorem 35 below is the main result in this section.

THEOREM 35. *For any $\alpha > 0$, any $k \geq 1$, and any $r \geq 1$, there exists a matrix $\mathbf{A} \in \mathbb{R}^{(n+1) \times n}$ for which*

$$\frac{\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_2^2}{\|\mathbf{A} - \mathbf{A}_k\|_2^2} \geq \frac{n + \alpha^2}{r + \alpha^2}.$$

Here \mathbf{C} is any matrix that consists of r columns of \mathbf{A} . As $\alpha \rightarrow 0$, the lower bound is n/r for the approximation ratio of spectral norm column-based matrix reconstruction.

Proof. We extend the lower bound in [6] to arbitrary $r > k$. Consider the matrix

$$\mathbf{A} = [\mathbf{e}_1 + \alpha\mathbf{e}_2, \mathbf{e}_1 + \alpha\mathbf{e}_3, \dots, \mathbf{e}_1 + \alpha\mathbf{e}_{n+1}] \in \mathbb{R}^{(n+1) \times n},$$

where $\mathbf{e}_i \in \mathbb{R}^{n+1}$ are the standard basis vectors. Then

$$\mathbf{A}^T\mathbf{A} = \mathbf{1}_n\mathbf{1}_n^T + \alpha^2\mathbf{I}_n, \quad \sigma_1^2(\mathbf{A}) = n + \alpha^2, \quad \text{and} \quad \sigma_i^2(\mathbf{A}) = \alpha^2 \quad \text{for } i > 1.$$

Thus, for all $k \geq 1$, $\|\mathbf{A} - \mathbf{A}_k\|_2^2 = \alpha^2$. Intuitively, as $\alpha \rightarrow 0$, \mathbf{A} is a rank-one matrix. Consider any r columns of \mathbf{A} and note that, up to row permutations, all sets of r columns of \mathbf{A} are equivalent. So, without loss of generality, let \mathbf{C} consist of the first r columns of \mathbf{A} . We now compute the optimal reconstruction of \mathbf{A} from \mathbf{C} as follows: let \mathbf{a}_j be the j th column of \mathbf{A} . In order to reconstruct \mathbf{a}_j , we minimize $\|\mathbf{a}_j - \mathbf{C}\mathbf{x}\|_2^2$ over all vectors $\mathbf{x} \in \mathbb{R}^r$. Note that if $j \leq r$, then the reconstruction error is zero. For $j > r$, $\mathbf{a}_j = \mathbf{e}_1 + \alpha\mathbf{e}_{j+1}$,

$$\mathbf{C}\mathbf{x} = \mathbf{e}_1 \sum_{i=1}^r x_i + \alpha \sum_{i=1}^r x_i \mathbf{e}_{i+1}.$$

Then

$$\begin{aligned} \|\mathbf{a}_j - \mathbf{C}\mathbf{x}\|_2^2 &= \|\mathbf{e}_1 \left(\sum_{i=1}^r x_i - 1 \right) + \alpha \sum_{i=1}^r x_i \mathbf{e}_{i+1} - \mathbf{e}_{j+1}\|_2^2 \\ &= \left(\sum_{i=1}^r x_i - 1 \right)^2 + \alpha^2 \sum_{i=1}^r x_i^2 + 1. \end{aligned}$$

The above quadratic form in \mathbf{x} is minimized when $x_i = (r + \alpha^2)^{-1}$ for all $i = 1, \dots, r$. Let $\hat{\mathbf{A}} = \mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}$ and let the j th column of $\hat{\mathbf{A}}$ be $\hat{\mathbf{a}}_j$. Then, for $j \leq r$, $\hat{\mathbf{a}}_j$ is an all-zeros vector; for $j > r$, $\hat{\mathbf{a}}_j = \alpha\mathbf{e}_{j+1} - \frac{\alpha}{r + \alpha^2} \sum_{i=1}^r \mathbf{e}_{i+1}$. Thus,

$$\hat{\mathbf{A}}^T \hat{\mathbf{A}} = \begin{bmatrix} \mathbf{0}_{r \times r} & \mathbf{0}_{r \times (n-r)} \\ \mathbf{0}_{(n-r) \times r} & \mathbf{Z} \end{bmatrix},$$

where

$$\mathbf{Z} = \frac{\alpha^2}{r + \alpha^2} \mathbf{1}_{n-r} \mathbf{1}_{n-r}^T + \alpha^2 \mathbf{I}_{n-r}.$$

This immediately implies that

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_2^2 = \|\hat{\mathbf{A}}\|_2^2 = \|\hat{\mathbf{A}}^T \hat{\mathbf{A}}\|_2 = \|\mathbf{Z}\|_2^2 = \frac{(n-r)\alpha^2}{r + \alpha^2} + \alpha^2 = \frac{n + \alpha^2}{r + \alpha^2} \alpha^2.$$

This concludes our proof, because $\alpha^2 = \|\mathbf{A} - \mathbf{A}_k\|_2^2$. \square

9.2. Frobenius norm approximation. Note that a lower bound for the ratio

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\xi}(\mathbf{A})\|_{\xi}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\xi}^2$$

does not imply a lower bound for the ratio

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\xi}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\xi}^2,$$

because

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\xi}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\xi}^2 \leq \|\mathbf{A} - \Pi_{\mathbf{C},k}^{\xi}(\mathbf{A})\|_{\xi}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\xi}^2.$$

Also, note that Proposition 4 in [9] shows a lower bound equal to $(1 + k/2r)$ for the ratio

$$\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbb{F}}(\mathbf{A})\|_{\mathbb{F}}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2.$$

For completeness, we extend the bound of [9] to the ratio

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbb{F}}^2 / \|\mathbf{A} - \mathbf{A}_k\|_{\mathbb{F}}^2.$$

In fact, we obtain a lower bound which is asymptotically $1 + k/r$. We start with the following lemma.

LEMMA 36. For any $\alpha > 0$ and $r \geq 1$, there exists a matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ for which

$$\frac{\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbb{F}}^2}{\|\mathbf{A} - \mathbf{A}_1\|_{\mathbb{F}}^2} \geq \frac{n-r}{n-1} \left(1 + \frac{1}{r + \alpha^2}\right).$$

Proof. We use the same matrix $\mathbf{A} \in \mathbb{R}^{(n+1) \times n}$ from Theorem 35. So, it follows that $\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbb{F}}^2 = \text{Tr}(\mathbf{Z}) = \alpha^2(n-r)(1 + \frac{1}{r+\alpha^2})$, and $\|\mathbf{A} - \mathbf{A}_1\|_{\mathbb{F}}^2 = (n-1)\alpha^2$, which concludes the proof. \square

THEOREM 37. For any $\alpha > 0$, any $k \geq 1$, and any $r \geq 1$, there exists a matrix $\mathbf{B} \in \mathbb{R}^{m \times \ell}$ for which

$$\frac{\|\mathbf{B} - \mathbf{C}\mathbf{C}^+\mathbf{B}\|_{\mathbb{F}}^2}{\|\mathbf{B} - \mathbf{B}_k\|_{\mathbb{F}}^2} \geq \frac{\ell-r}{\ell-k} \left(1 + \frac{k}{r + \alpha^2}\right).$$

Here \mathbf{C} is any matrix that consists of r columns of \mathbf{B} . By taking $\alpha \rightarrow 0$ and $\ell \rightarrow \infty$ the lower bound is $1 + (k/r)$ for the approximation ratio of Frobenius norm column-based matrix reconstruction.

In our construction of \mathbf{B} , $m = (n+1)k$ and $\ell = nk$ for any $\ell \geq r$.

Proof. The matrix \mathbf{B} is constructed as follows. Let \mathbf{A} have dimensions $(n+1) \times n$ and be constructed as in Theorem 35, except with α replaced by $\alpha' = \alpha/\sqrt{k}$. \mathbf{B} is block diagonal, with k copies of \mathbf{A} along its diagonal; so, the dimensions of \mathbf{B} are $m = (n+1)k$ and $\ell = nk$. We sample r columns in total, with r_i columns from each block. Lemma 36 holds for each block, with r replaced by r_i , where $r_i \geq 0$ and $\sum_{i=1}^k r_i = r$. We analyze the Frobenius norm error in each block independently. Let \mathbf{Z}_i be the error matrix in each block, as in the proof of Theorem 35. Then, using Lemma 36, the approximation error is equal to

$$\|\mathbf{B} - \mathbf{C}\mathbf{C}^+\mathbf{B}\|_{\mathbb{F}}^2 = \sum_{i=1}^k \text{Tr}(\mathbf{Z}_i) = \alpha'^2 \sum_{i=1}^k (n-r_i) \left(1 + \frac{1}{r_i + \alpha'^2}\right).$$

This last expression is minimized subject to the constraint that $\sum_{i=1}^k r_i = r$ when $r_i = r/k$, and so

$$\|\mathbf{B} - \mathbf{C}\mathbf{C}^+\mathbf{B}\|_{\mathbb{F}}^2 \geq \alpha'^2(nk-r) \left(1 + \frac{k}{r + k\alpha'^2}\right) = \alpha'^2(\ell-r) \left(1 + \frac{k}{r + \alpha^2}\right),$$

where we used $\alpha^2 = k\alpha'^2$. The result follows because

$$\|\mathbf{B} - \mathbf{B}_k\|_{\mathbb{F}}^2 = k\|\mathbf{A} - \mathbf{A}_1\|_{\mathbb{F}}^2 = k(n-1)\alpha'^2 = (\ell-k)\alpha'^2. \quad \square$$

10. Conclusions and open problems. Several interesting questions remain unanswered, which we summarize in this section.

First, is it possible to improve the running time of the deterministic algorithms of Lemmas 13 and 14? Recently, Zouzias [26] made progress in improving the running time of the spectral sparsification result of [2]; can we get a similar improvement for the 2-set algorithms presented here? Or perhaps, can we trade off the running time with randomization in those algorithms?

Second, in the parlance of Theorem 5, is there a *deterministic* algorithm that selects $O(k/\epsilon)$ columns from \mathbf{A} and guarantees relative-error accuracy for the error $\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbf{F}}(\mathbf{A})\|_{\mathbf{F}}^2$? Such a deterministic bound would be possible, for example, by derandomizing the adaptive sampling technique used in Theorem 5. In a recent development, [18] partially answers this question by extending the volume sampling approach of [6] to deterministically select $\frac{k}{\epsilon}(1 + o(1))$ columns and obtain a relative error bound for the term $\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbf{F}}^2$. Notice that it is not obvious if [18] implies a similar deterministic bound for the error $\|\mathbf{A} - \Pi_{\mathbf{C},k}^{\mathbf{F}}(\mathbf{A})\|_{\mathbf{F}}^2$.

Third, is it possible to develop a (deterministic or randomized) algorithm to compute the best (in the *spectral* norm) rank restricted projection of \mathbf{A} into a specific subspace \mathbf{C} efficiently? That is, can one compute $\Pi_{\mathbf{C},k}^2(\mathbf{A})$ in time that is $O(T_{SVD}(\mathbf{C}))$?

Finally, is it possible to develop column reconstruction algorithms that provide upper bounds in terms of the column reconstruction that is achieved with an optimal set of columns? For example, given \mathbf{A} , k , and $r \geq k$, let \mathbf{C}_{opt} be the best choice of r columns in \mathbf{A} . Is it possible to develop algorithms that select matrices \mathbf{C} with r columns and provide upper bounds of the form

$$\|\mathbf{A} - \mathbf{C}\mathbf{C}^+\mathbf{A}\|_{\mathbf{F}}^2 \leq \gamma \cdot \|\mathbf{A} - \mathbf{C}_{opt}\mathbf{C}_{opt}^+\mathbf{A}\|_{\mathbf{F}}^2,$$

where $\gamma \geq 1$ is the approximation factor? Also, how do the lower bounds change if we replace $\mathbf{A} - \mathbf{A}_k$ with $\mathbf{A} - \mathbf{C}_{opt}\mathbf{C}_{opt}^+\mathbf{A}$?

Acknowledgments. We would like to thank A. Deshpande, D. Feldman, K. Varadarajan, and J. Tropp for useful discussions, and D. Feldman and K. Varadarajan for pointing out the connections between the subspace approximation line of research [8, 13, 14, 24] and our work. We would also like to thank an anonymous reviewer for pointing out [25] and how this result on oblique projections implies (2) and its implications to Theorems 3 and 15; this avenue suggested by the reviewer slightly improved our previous bounds.

REFERENCES

- [1] H. AVRON AND C. BOUTSIDIS, *Faster subset selection for matrices and applications*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1464–1499.
- [2] J.D. BATSON, D.A. SPIELMAN, AND N. SRIVASTAVA, *Twice-Ramanujan sparsifiers*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), 2009, pp. 255–262.
- [3] C. BOUTSIDIS, M.W. MAHONEY, AND P. DRINEAS, *An improved approximation algorithm for the column subset selection problem*, in Proceedings of the 20th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2009, pp. 968–977.
- [4] T.F. CHAN AND P.C. HANSEN, *Some applications of the rank revealing QR factorization*, SIAM J. Sci. Statist. Comput., 13 (1992), pp. 727–741.
- [5] K.L. CLARKSON AND D.P. WOODRUFF, *Numerical linear algebra in the streaming model*, in Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC), 2009.
- [6] A. DESHPANDE AND L. RADEMACHER, *Efficient volume sampling for row/column subset selection*, in Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC), 2010.

- [7] A. DESHPANDE, L. RADEMACHER, S. VEMPALA, AND G. WANG, *Matrix approximation and projective clustering via volume sampling*, Theory Comput., 2 (2006), pp. 225–247.
- [8] A. DESHPANDE AND K.R. VARADARAJAN, *Sampling-based dimension reduction for subspace approximation*, in Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC), 2007.
- [9] A. DESHPANDE AND S. VEMPALA, *Adaptive sampling and fast low-rank matrix approximation*, in Proceedings of RANDOM - APPROX, 2006.
- [10] P. DRINEAS, A. FRIEZE, R. KANNAN, S. VEMPALA, AND V. VINAY, *Clustering in large graphs and matrices*, in Proceedings of the 10th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 1999, pp. 291–299.
- [11] P. DRINEAS, I. KERENIDIS, AND P. RAGHAVAN, *Competitive recommendation systems*, in Proceedings of 34th Annual ACM Symposium on Theory of Computing (STOC), 2002.
- [12] P. DRINEAS, M.W. MAHONEY, AND S. MUTHUKRISHNAN, *Polynomial Time Algorithm for Column-Row Based Relative-Error Low-Rank Matrix Approximation*, Report, DIMACS, 2006.
- [13] D. FELDMAN AND M. LANGBERG, *A unified framework for approximating and clustering data*, in Proceedings of the 43rd Annual ACM Symposium on Theory of Computing (STOC), 2011.
- [14] D. FELDMAN, M. MONEMIZADEH, C. SOHLER, AND D.P. WOODRUFF, *Coresets and sketches for high dimensional subspace approximation problems*, in Proceedings of the 21st Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2010, pp. 630–649.
- [15] A. FRIEZE, R. KANNAN, AND S. VEMPALA, *Fast Monte-Carlo algorithms for finding low-rank approximations*, in Proceedings from the 39th Annual Symposium on Foundations of Computer Science (FOCS), 1998.
- [16] G. GOLUB, *Numerical methods for solving linear least squares problems*, Numer. Math., 7 (1965), pp. 206–216.
- [17] M. GU AND S.C. EISENSTAT, *Efficient algorithms for computing a strong rank-revealing QR factorization*, SIAM J. Sci. Comput., 17 (1996), pp. 848–869.
- [18] V. GURUSWAMI AND A.K. SINOP, *Optimal column-based low-rank matrix reconstruction*, in Proceedings of the 23rd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2012, pp. 1207–1214.
- [19] N. HALKO, P.G. MARTINSSON, AND J.A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Rev., 53 (2011), pp. 217–288.
- [20] E. LIBERTY, F. WOOLFE, P.G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *Randomized algorithms for the low-rank approximation of matrices*, Proc. Natl. Acad. Sci. USA, 104 (2007), pp. 20167–20172.
- [21] M.W. MAHONEY AND P. DRINEAS, *CUR matrix decompositions for improved data analysis*, Proc. Natl. Acad. Sci. USA, 106 (2009), pp. 697–702.
- [22] V. ROKHLIN, A. SZLAM, AND M. TYGERT, *A randomized algorithm for principal component analysis*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1100–1124.
- [23] T. SARLOS, *Improved approximation algorithms for large matrices via random projections*, in Proceedings of the 47th Annual Symposium on Foundations of Computer Science (FOCS), 2006.
- [24] N.D. SHYAMALKUMAR AND K. VARADARAJAN, *Efficient subspace approximation algorithms*, in Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2007, pp. 532–540.
- [25] D. SZYLD, *The many proofs of an identity on the norm of oblique projections*, Numer. Algorithms, 42 (2006), pp. 309–323.
- [26] A. ZOUZIAS, *A matrix hyperbolic cosine algorithm and applications*, in Proceedings of the 39th International Colloquium on Automata, Languages and Programming (ICALP), 2012.