COLUMN SUBSET SELECTION, MATRIX FACTORIZATION, AND EIGENVALUE OPTIMIZATION

JOEL A. TROPP

Technical Report No. 2008-02
July 2008

APPLIED & COMPUTATIONAL MATHEMATICS
CALIFORNIA INSTITUTE OF TECHNOLOGY
mail code 217-50 · pasadena, ca 91125

COLUMN SUBSET SELECTION, MATRIX FACTORIZATION, AND EIGENVALUE OPTIMIZATION

J. A. TROPP

Abstract. Given a fixed matrix, the problem of column subset selection requests a column submatrix that has favorable spectral properties. Most research from the algorithms and numerical linear algebra communities focuses on a variant called rank-revealing QR, which seeks a well-conditioned collection of columns that spans the (numerical) range of the matrix. The functional analysis literature contains another strand of work on column selection whose algorithmic implications have not been explored. In particular, a celebrated result of Bourgain and Tzafriri demonstrates that each matrix with normalized columns contains a large column submatrix that is exceptionally well conditioned. Unfortunately, standard proofs of this result cannot be regarded as algorithmic.

This paper presents a randomized, polynomial-time algorithm that produces the submatrix promised by Bourgain and Tzafriri. The method involves random sampling of columns, followed by a matrix factorization that exposes the well-conditioned subset of columns. This factorization, which is due to Grothendieck, is regarded as a central tool in modern functional analysis. The primary novelty in this work is an algorithm, based on eigenvalue minimization, for constructing the Grothendieck factorization. These ideas also result in a novel approximation algorithm for the $(\infty, 1)$ norm of a matrix, which is generally NP-hard to compute exactly. As an added bonus, this work reveals a surprising connection between matrix factorization and the famous MAXCUT semidefinite program.

Date: 26 June 2008.

JAT is with Applied and Computational Mathematics, MC 217-50, California Inst. Technology, Pasadena, CA 91125-5000. E-mail: jtropp@acm.caltech.edu. Supported in part by ONR award no. N000140810883.

1. Introduction

Column subset selection refers to the challenge of extracting from a matrix a column submatrix that has some distinguished property. These properties commonly involve conditions on the spectrum of the submatrix. The most familiar example is probably rank-revealing QR, which seeks a well-conditioned collection of columns that spans the (numerical) range of the matrix [GE96].

The literature on geometric functional analysis contains several fundamental theorems on column subset selection that have not been discussed by the algorithms community or the numerical linear algebra community. These results are phrased in terms of the *stable rank* of a matrix:

$$\operatorname{st.rank}(\boldsymbol{A}) = \frac{\|\boldsymbol{A}\|_{\mathrm{F}}^2}{\|\boldsymbol{A}\|^2}$$

where $\|\cdot\|_{\mathrm{F}}$ is the Frobenius norm and $\|\cdot\|$ is the spectral norm. The stable rank can be viewed as an analytic surrogate for the algebraic rank. Indeed, express the two norms in terms of singular values to obtain the relation

$$\operatorname{st.rank}(\boldsymbol{A}) \le \operatorname{rank}(\boldsymbol{A}).$$

In this bound, equality occurs (for example) when the columns of $\boldsymbol{A}$ are identical or when the columns of $\boldsymbol{A}$ are orthonormal. As we will see, the stable rank is tightly connected with the number of (strongly) linearly independent columns we can extract from a matrix.

Before we continue, let us instate some notation. We say that a matrix is *standardized* when its columns have unit ℓ_2 norm. The jth column of a matrix $\boldsymbol{A}$ is denoted by $\boldsymbol{a}_j$. For a subset τ of column indices, we write $\boldsymbol{A}_\tau$ for the column submatrix indexed by τ. Likewise, given a square matrix $\boldsymbol{H}$, the notation $\boldsymbol{H}_{\tau\times\tau}$ refers to the principal submatrix whose rows and columns are listed in τ. The pseudoinverse $\boldsymbol{D}^\dagger$ of a diagonal matrix $\boldsymbol{D}$ is formed by reciprocating the nonzero entries. As usual, we write $\|\cdot\|_p$ for the ℓ_p vector norm. The *condition number* of a matrix is the quantity

$$\kappa(\boldsymbol{A}) = \max\left\{ \frac{\|\boldsymbol{A}\boldsymbol{x}\|_2}{\|\boldsymbol{A}\boldsymbol{y}\|_2} : \|\boldsymbol{x}\|_2 = \|\boldsymbol{y}\|_2 = 1 \right\}.$$

Finally, upright letters $(\mathrm{c}, \mathrm{C}, \mathrm{K}, \dots)$ refer to positive, universal constants that may change from appearance to appearance.

The first theorem, due to Kashin and Tzafriri, shows that each matrix with standardized columns contains a large column submatrix that has small spectral norm [Ver01, Thm. 2.5].

Theorem 1.1 (Kashin–Tzafriri)**.** *Suppose $\boldsymbol{A}$ is standardized. Then there is a set τ of column indices for which*

$$|\tau| \ge \operatorname{st.rank}(\boldsymbol{A}) \quad \textit{and} \quad \|\boldsymbol{A}_\tau\| \le \mathrm{C}.$$

In fact, much more is true. Combining Theorem 1.1 with the celebrated restricted invertibility result of Bourgain and Tzafriri [BT87, Thm. 1.2], we find that every standardized matrix contains a large column submatrix whose *condition number* is small.

Theorem 1.2 (Bougain–Tzafriri)**.** *Suppose $\boldsymbol{A}$ is standardized. Then there is a set τ of column indices for which*

$$|\tau| \ge \mathrm{c} \cdot \operatorname{st.rank}(\boldsymbol{A}) \quad \textit{and} \quad \kappa(\boldsymbol{A}_\tau) \le \sqrt{3}.$$

Theorem 1.2 yields the best general result [BT91, Thm. 1.1] on the Kadison–Singer conjecture, the major open question in operator theory. To display its strength, let us consider two extreme examples.

(1) When $\boldsymbol{A}$ has identical columns, every collection of two or more columns is singular. Theorem 1.2 guarantees a well-conditioned submatrix $\boldsymbol{A}_\tau$ with $|\tau| = 1$, which is optimal.
(2) When $\boldsymbol{A}$ has n orthonormal columns, the full matrix is perfectly conditioned. Theorem 1.2 guarantees a well-conditioned submatrix $\boldsymbol{A}_\tau$ with $|\tau| \ge \mathrm{c}n$, which lies within a constant factor of optimal.

The stable rank allows Theorem 1.2 to interpolate between the two extremes. Subsequent research established that the stable rank is intrinsic to the problem of finding well-conditioned submatrices. We postpone a more detailed discussion of this point until Section 6.

1.1. **Contributions.** Although Theorems 1.1 and 1.2 would be very useful in computational applications, we cannot regard current proofs as constructive. The goal of this paper is to establish the following novel, algorithmic claim.

Theorem 1.3. *There are randomized, polynomial-time algorithms for producing the sets guaranteed by Theorem 1.1 and by Theorem 1.2.*

This result is significant because no known algorithm for column subset selection is guaranteed to produce a submatrix whose condition number has constant order. See [BDM08] for a recent overview of that literature. The present work has other ramifications with independent interest.

- We develop algorithms for computing the matrix factorizations of Pietsch and Grothendieck, which are regarded as basic instruments in modern functional analysis [Pis86].
- The methods for computing these factorizations lead to new approximation algorithms for two NP-hard matrix norms. (See Remarks 3.2 and 5.6.)
- We identify an intriguing connection between Pietsch factorization and the MAXCUT semidefinite program [GW95].

1.2. **Overview.** We focus on the algorithmic version of the Kashin–Tzafriri theorem because it highlights all the essential concepts while minimizing irrelevant details. Section 2 outlines a proof of this result, emphasizing where new algorithmic machinery is required. The missing link turns out to be a computational method for producing a certain matrix factorization. Section 3 reformulates the factorization problem as an eigenvalue minimization, which can be completed with standard techniques. In Section 4, we exhibit a randomized algorithm that delivers the submatrix promised by Kashin–Tzafriri. In Section 5, we traverse a similar route to develop an algorithmic version of Bourgain–Tzafriri. Section 6 provides more details about the stable rank and describes directions for future work. Appendix A contains some key estimates on the norms of random submatrices, and Appendix B outlines a simple computational procedure for solving the eigenvalue optimization problems that arise in our work.

2. The Kashin–Tzafriri Theorem

The proof of the Kashin–Tzafriri theorem proceeds in two steps. First, we select a random set of columns with appropriate cardinality. Second, we use a matrix factorization to identify and remove redundant columns that inflate the spectral norm. The proof gives strong hints about how a computational procedure might work, even though it is not constructive.

2.1. **Intuitions.** We would like to think that a random submatrix inherits its share of the norm of the entire matrix. In other words, if we were to select a tenth of the columns, we might hope to reduce the norm by a factor of ten. Unfortunately, this intuition is meretricious.

Indeed, random selection does not necessarily reduce the spectral norm at all. The essential reason emerges when we consider the "double identity," the $m \times 2m$ matrix $\boldsymbol{A} = \begin{bmatrix}\mathbf{I} \mid \mathbf{I}\end{bmatrix}$. Suppose we draw s random columns from $\boldsymbol{A}$ without replacement. The probability that all s columns are distinct is

$$\frac{2m-2}{2m-1} \times \frac{2m-4}{2m-2} \times \cdots \times \frac{2m-2(s-1)}{2m-(s-1)} \leq \prod\nolimits_{j=0}^{s-1} \left(1 - \frac{j}{2m}\right) \approx \exp\left\{-\sum\nolimits_{j=0}^{s-1} \frac{j}{2m}\right\} \approx \mathrm{e}^{-s^2/4m}.$$

Therefore, when $s = \Omega(\sqrt{m})$, sampling almost always produces a submatrix with at least one duplicated column. A duplicated column means that the norm of the submatrix is $\sqrt{2}$, which equals the norm of the full matrix, so no reduction takes place.

Nevertheless, a randomly chosen set of columns from a standardized matrix typically *contains* a large set of columns that has small norm. We will see that the desired subset is exposed by factoring the random submatrix. This factorization, which was invented by Pietsch, is regarded as a basic instrument in modern functional analysis.

2.2. **The $(\infty, 2)$ operator norm.** Although sampling does not necessarily reduce the spectral norm, it often reduces other matrix norms. Define the natural norm on linear operators from ℓ_∞ to ℓ_2 via the expression

$$\|\boldsymbol{B}\|_{\infty\to 2} = \max\{\|\boldsymbol{B}\boldsymbol{x}\|_2 : \|\boldsymbol{x}\|_\infty = 1\}.$$

An immediate consequence is that $\|\boldsymbol{B}\|_{\infty\to 2} \le \sqrt{s}\,\|\boldsymbol{B}\|$ for each matrix $\boldsymbol{B}$ with s columns. Equality can obtain in this bound.

The exact calculation of the $(\infty, 2)$ operator norm is computationally difficult. Results of Rohn [Roh00] imply that there is a class of positive semidefinite matrices for which it is NP-hard to estimate $\|\cdot\|_{\infty\to 2}$ within an absolute tolerance. Nevertheless, we will see that the norm can be approximated in polynomial time up to a small relative error. (See Remark 3.2.)

As we have intimated, the $(\infty, 2)$ norm can often be reduced by random selection. The following theorem requires some heavy lifting, which we delegate to Appendix A.2.

Theorem 2.1. *Suppose $\boldsymbol{A}$ is a standardized matrix with n columns. Choose*

$$s \le \lceil 2\,\mathrm{st.rank}(\boldsymbol{A}) \rceil,$$

and draw a uniformly random subset σ with cardinality s from $\{1, 2, \dots, n\}$. Then

$$\mathbb{E}\,\|\boldsymbol{A}_\sigma\|_{\infty\to 2} \le 7\sqrt{s}.$$

In particular, $\|\boldsymbol{A}_\sigma\|_{\infty\to 2} \le 8\sqrt{s}$ with probability at least $1/8$.

2.3. **Pietsch factorization.** We cannot exploit the bound in Theorem 2.1 unless we have a way to connect the $(\infty, 2)$ norm with the spectral norm. To that end, let us recall one of the landmark theorems of functional analysis.

Theorem 2.2 (Pietsch Factorization)**.** *Each matrix $\boldsymbol{B}$ can be factored as $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ where*

- *$\boldsymbol{D}$ is a nonnegative, diagonal matrix with $\mathrm{trace}(\boldsymbol{D}^2) = 1$, and*
- *$\|\boldsymbol{B}\|_{\infty\to 2} \le \|\boldsymbol{T}\| \le \mathrm{K_P}\,\|\boldsymbol{B}\|_{\infty\to 2}$.*

This result follows from the little Grothendieck theorem [Pis86, Sec. 5b] and the Pietsch factorization theorem [Pis86, Cor. 1.8]. The standard proof produces the factorization using an abstract separation argument that offers no algorithmic insight. The value of the constant is available.

- When the scalar field is real, we have $\mathrm{K_P}(\mathbb{R}) = \sqrt{\pi/2} \approx 1.25$.
- When the scalar field is complex, we have $\mathrm{K_P}(\mathbb{C}) = \sqrt{4/\pi} \approx 1.13$.

A major application of Pietsch factorization is to identify a submatrix with controlled spectral norm. The following proposition describes the procedure.

Proposition 2.3. *Suppose $\boldsymbol{B}$ is a matrix with s columns. Then there is a set τ of column indices for which*

$$|\tau| \ge \frac{s}{2} \quad \text{and} \quad \|\boldsymbol{B}_\tau\| \le \mathrm{K_P}\sqrt{\frac{2}{s}}\,\|\boldsymbol{B}\|_{\infty\to 2}\,.$$

Proof. Consider a Pietsch factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$, and define

$$\tau = \{j : d_{jj}^2 \le 2/s\}.$$

Since $\sum d_{jj}^2 = 1$, Markov's inequality implies that $|\tau| \ge s/2$. We may calculate that

$$\|\boldsymbol{B}_\tau\| = \|\boldsymbol{T}\boldsymbol{D}_\tau\| \le \|\boldsymbol{T}\| \cdot \|\boldsymbol{D}_\tau\| \le \mathrm{K_P}\,\|\boldsymbol{B}\|_{\infty\to 2} \cdot \sqrt{2/s}.$$

This completes the proof. □

2.4. **Proof of Kashin–Tzafriri.** With these results at hand, we easily complete the proof of the Kashin–Tzafriri theorem. Suppose $\boldsymbol{A}$ is a standardized matrix with n columns. Assume that $\operatorname{st.rank}(\boldsymbol{A}) \le n/2$. Otherwise, the spectral norm $\|\boldsymbol{A}\| \le \sqrt{2}$, so we may select $\tau = \{1, 2, \dots, n\}$.

According to Theorem 2.1, there is a subset σ of column indices for which

$$|\sigma| \ge 2 \operatorname{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \|\boldsymbol{A}_\sigma\|_{\infty\to 2} \le 8\sqrt{|\sigma|}.$$

Apply Proposition 2.3 to the matrix $\boldsymbol{B} = \boldsymbol{A}_\sigma$ to obtain a subset τ inside σ for which

$$|\tau| \ge \frac{|\sigma|}{2} \quad \text{and} \quad \|\boldsymbol{B}_\tau\| \le \mathrm{K_P} \sqrt{\frac{2}{|\sigma|}} \, \|\boldsymbol{B}\|_{\infty\to 2}\,.$$

Since $\boldsymbol{B}_\tau = \boldsymbol{A}_\tau$ and $\mathrm{K_P} \le \sqrt{\pi/2}$, these bounds reveal the advertised conclusion:

$$|\tau| \ge \operatorname{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \|\boldsymbol{A}_\tau\| < 15.$$

At this point, we take a step back and notice that this proof is nearly algorithmic. It is straightforward to perform the random selection described in Theorem 2.1. Provided that we know a Pietsch factorization of the matrix $\boldsymbol{B}$, we can easily carry out the column selection of Proposition 2.3. Therefore, we need only develop an algorithm for computing the Pietsch factorization to reach an effective version of the Kashin–Tzafriri theorem.

3. Pietsch Factorization via Convex Optimization

The main novelty is to demonstrate that we can produce a Pietsch factorization by solving a convex programming problem. Remarkably, the resulting optimization is the dual of the famous MAXCUT semidefinite program [GW95], for which many polynomial-time algorithms are available.

3.1. **Pietsch and eigenvalues.** The next theorem, which serves as the basis for our computational method, demonstrates that Pietsch factorizations have an intimate relationship with the eigenvalues of a related matrix. In the sequel, we reserve the letter $\boldsymbol{D}$ for a nonnegative, diagonal matrix with $\operatorname{trace}(\boldsymbol{D}^2) = 1$, and we write $\lambda_{\max}$ for the algebraically maximal eigenvalue of a Hermitian matrix.

Theorem 3.1. *The factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ satisfies $\|\boldsymbol{T}\| \le \alpha$ if and only if $\boldsymbol{D}$ satisfies*

$$\lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2 \boldsymbol{D}^2) \le 0.$$

In particular, if no $\boldsymbol{D}$ verifies this bound, then no factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ admits $\|\boldsymbol{T}\| \le \alpha$.

Proof. Assume $\boldsymbol{B}$ has a factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ with $\|\boldsymbol{T}\| \le \alpha$. We have the chain of implications

$$\begin{aligned}
\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D} \quad &\implies \quad \|\boldsymbol{B}\boldsymbol{x}\|_2^2 = \|\boldsymbol{T}\boldsymbol{D}\boldsymbol{x}\|_2^2 && \forall \boldsymbol{x} \\
&\implies \quad \|\boldsymbol{B}\boldsymbol{x}\|_2^2 \le \alpha^2 \, \|\boldsymbol{D}\boldsymbol{x}\|_2^2 && \forall \boldsymbol{x} \\
&\implies \quad \boldsymbol{x}^*\boldsymbol{B}^*\boldsymbol{B}\boldsymbol{x} \le \alpha^2 \boldsymbol{x}^*\boldsymbol{D}^2\boldsymbol{x} && \forall \boldsymbol{x} \\
&\implies \quad \boldsymbol{x}^*(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{D}^2)\boldsymbol{x} \le 0 && \forall \boldsymbol{x} \\
&\implies \quad \boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{D}^2 \preccurlyeq \mathbf{0},
\end{aligned}$$

where $\preccurlyeq$ denotes the semidefinite, or Löwner, ordering on Hermitian matrices.

Conversely, assume we are provided the inequality

$$\boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{D}^2 \preccurlyeq \mathbf{0}. \tag{3.1}$$

First, we claim that any zero entry in $\boldsymbol{D}$ corresponds with a zero column of $\boldsymbol{B}$. To check this point, suppose that $d_{jj} = 0$ for an index j. The relation (3.1) requires that

$$0 \ge (\boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{D}^2)_{jj} = \boldsymbol{b}_j^*\boldsymbol{b}_j.$$

This inequality is impossible unless $\boldsymbol{b}_j = \mathbf{0}$. To continue, set $\boldsymbol{T} = \boldsymbol{B}\boldsymbol{D}^\dagger$, and observe that $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ because the zero entries of $\boldsymbol{D}$ correspond with zero columns of $\boldsymbol{B}$. Therefore, we may factor the diagonal matrix out from (3.1) to reach

$$\boldsymbol{D}(\boldsymbol{T}^*\boldsymbol{T} - \alpha^2\boldsymbol{P})\boldsymbol{D} \preccurlyeq \mathbf{0}.$$

where the matrix $\boldsymbol{P} = \boldsymbol{D}\boldsymbol{D}^\dagger$ is an orthogonal projector. Sylvester's theorem on inertia [HJ85, Thm. 4.5.8] ensures that $\boldsymbol{T}^*\boldsymbol{T} - \alpha^2\boldsymbol{P} \preccurlyeq \mathbf{0}$. Since $\boldsymbol{P}$ is a projector, this relation implies that

$$\boldsymbol{T}^*\boldsymbol{T} \preccurlyeq \alpha^2\boldsymbol{P} \preccurlyeq \alpha^2\mathbf{I}.$$

We conclude that $\|\boldsymbol{T}\| \le \alpha$. □

3.2. **Factorization via optimization.** Recall that the maximum eigenvalue is a convex function on the space of Hermitian matrices, so it can be minimized in polynomial time [LO96]. We are led to consider the convex program

$$\min\ \lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{F}) \qquad \text{subject to} \qquad \operatorname{trace}(\boldsymbol{F}) = 1,\ \boldsymbol{F} \text{ diagonal, and } \boldsymbol{F} \ge \mathbf{0}. \tag{3.2}$$

Owing to Theorem 3.1, there exists a factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$ with $\|\boldsymbol{T}\| \le \alpha$ if and only if the value of (3.2) is nonpositive.

Now, if $\boldsymbol{F}$ is a feasible point of (3.2) with a nonpositive objective value, we can factorize

$$\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D} \quad \text{with} \quad \boldsymbol{D} = \boldsymbol{F}^{1/2}, \quad \boldsymbol{T} = \boldsymbol{B}\boldsymbol{D}^\dagger, \quad \text{and} \quad \|\boldsymbol{T}\| \le \alpha.$$

In fact, it is not necessary to solve (3.2) to optimality. Suppose $\boldsymbol{B}$ has s columns, and assume we have identified a feasible point $\boldsymbol{F}$ with a (positive) objective value η. That is,

$$\lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2\boldsymbol{F}) \le \eta.$$

Rearranging this relation, we reach

$$\lambda_{\max}\left[\boldsymbol{B}^*\boldsymbol{B} - (\alpha^2 + \eta s)\widetilde{\boldsymbol{F}}\right] \le 0 \quad \text{where} \quad \widetilde{\boldsymbol{F}} = \frac{1}{\alpha^2 + \eta s}(\alpha^2\boldsymbol{F} + \eta\mathbf{I}).$$

Since $\widetilde{\boldsymbol{F}}$ is positive and diagonal with $\operatorname{trace}(\widetilde{\boldsymbol{F}}) = 1$, we obtain the factorization

$$\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D} \quad \text{with} \quad \boldsymbol{D} = \widetilde{\boldsymbol{F}}^{1/2}, \quad \boldsymbol{T} = \boldsymbol{B}\boldsymbol{D}^{-1}, \quad \text{and} \quad \|\boldsymbol{T}\| \le \sqrt{\alpha^2 + \eta s}.$$

To select a target value for the parameter α, we look to the proof of the Kashin–Tzafriri theorem. If $\boldsymbol{B}$ has s columns, then $\alpha = 8\mathrm{K_P}\sqrt{s}$ is an appropriate choice. Furthermore, since the argument only uses the bound $\|\boldsymbol{T}\| = \mathrm{O}(\sqrt{s})$, it suffices to solve (3.2) with precision $\eta = \mathrm{O}(1)$.

3.3. **Other formulations.** In a general setting, a target value for α is not likely to be available. Let us exhibit an alternative formulation of (3.2) that avoids this inconvenience.

$$\min\ \lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \boldsymbol{E}) + \operatorname{trace}(\boldsymbol{E}) \qquad \text{subject to} \qquad \boldsymbol{E} \text{ diagonal, } \boldsymbol{E} \ge \mathbf{0}. \tag{3.3}$$

Suppose $\alpha_\star$ is the minimal value of $\|\boldsymbol{T}\|$ achievable in any Pietsch factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$. It can be shown that $\alpha_\star^2$ is the value of (3.3) and that each optimizer $\boldsymbol{E}_\star$ satisfies $\operatorname{trace}(\boldsymbol{E}_\star) = \alpha_\star^2$. As such, we can construct an optimal Pietsch factorization from a minimizer:

$$\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D} \quad \text{with} \quad \boldsymbol{D} = (\boldsymbol{E}_\star/\operatorname{trace}(\boldsymbol{E}_\star))^{1/2}, \quad \boldsymbol{T} = \boldsymbol{B}\boldsymbol{D}^\dagger, \quad \text{and} \quad \|\boldsymbol{T}\| = \alpha_\star.$$

The dual of (3.3) is the semidefinite program

$$\max\ \langle \boldsymbol{B}^*\boldsymbol{B},\ \boldsymbol{Z}\rangle \qquad \text{subject to} \qquad \operatorname{diag}(\boldsymbol{Z}) = \mathbf{I} \text{ and } \boldsymbol{Z} \succcurlyeq \mathbf{0}. \tag{3.4}$$

This is the famous MAXCUT semidefinite program [GW95]. We find an unexpected connection between Pietsch factorization and the problem of partitioning nodes of a graph.

Given a dual optimum, we can easily construct a primal optimum by means of the complementary slackness condition [Ali95, Thm. 2.10]. Indeed, each feasible optimal pair $(\boldsymbol{E}_\star, \boldsymbol{Z}_\star)$ satisfies $\boldsymbol{Z}_\star(\boldsymbol{B}^*\boldsymbol{B} - \boldsymbol{E}_\star) = \mathbf{0}$. Examining the diagonal elements of this matrix equation, we find that

$$\boldsymbol{E}_\star = \operatorname{diag}(\boldsymbol{E}_\star) = \operatorname{diag}(\boldsymbol{Z}\boldsymbol{E}_\star) = \operatorname{diag}(\boldsymbol{Z}_\star\boldsymbol{B}^*\boldsymbol{B})$$

owing to the constraint $\operatorname{diag}(\boldsymbol{Z}_\star) = \mathbf{I}$. Obtaining a dual optimum from a primal optimum, however, requires more ingenuity.

Remark 3.2. *According to Theorem 2.2 and the discussion here, the optimal value of* (3.3) *overestimates* $\|\boldsymbol{B}\|_{\infty\to 2}^2$ *by a multiplicative factor no greater than* $\mathrm{K_P}^2$. *As a result, the optimization problem* (3.3) *can be used to design an approximation algorithm for* $(\infty, 2)$ *norms.*

3.4. **Algorithmic aspects.** The purpose of this paper is not to rehash methods for solving a standard optimization problem, so we keep this discussion brief. It is easy to see that (3.2) can be framed as a (nonsmooth) convex optimization over the probability simplex. Appendix B outlines an elegant technique, called Entropic Mirror Descent [BT03], designed specifically for this class of problems. Although the EMD algorithm is (theoretically) not the most efficient approach to (3.2), preliminary experiments suggest that its empirical performance rivals more sophisticated techniques.

For a concrete time bound, we refer to Alizadeh's work on primal–dual potential reduction methods for semidefinite programming [Ali95]. When $\boldsymbol{B}$ has dimension $m \times s$, the cost of forming $\boldsymbol{B}^*\boldsymbol{B}$ is at most $\mathrm{O}(s^2 m)$. Then the cost of solving (3.4) is no more than $\widetilde{\mathrm{O}}(s^{3.5})$, where the tilde indicates that log-like factors are suppressed.

4. An Algorithm for Kashin–Tzafriri

At this point, we have amassed the matériel necessary to deploy an algorithm that constructs the set τ promised by the Kashin–Tzafriri theorem. The procedure appears on page 11 as Algorithm 1. The following result describes its performance.

Theorem 4.1. *Suppose* $\boldsymbol{A}$ *is an* $m \times n$ *standardized matrix. With probability at least* $4/5$, *Algorithm 1 produces a set* $\tau = \tau_\star$ *of column indices for which*

$$|\tau| \geq \frac{1}{2}\operatorname{st.rank}(\boldsymbol{A}) \quad \textit{and} \quad \|\boldsymbol{A}_\tau\| \leq 15.$$

The computational cost is bounded by $\widetilde{\mathrm{O}}(|\tau|^2\, m + |\tau|^{3.5})$.

Remarkably, Algorithm 1 is sublinear in the size of the matrix when $\operatorname{st.rank}(\boldsymbol{A}) = \mathrm{o}(n^{1/3.5})$. Better methods for solving (3.2) would strengthen this bound.

Proof. According to Section 2, the procedure Norm-Reduce has failure probability less than $7/8$ when $s \leq 2\operatorname{st.rank}(\boldsymbol{A})$. The probability the inner loop fails to produce an acceptable set $\tau_\star$ of size $s/2$ is at most $(7/8)^{8\log_2(s)}$. So the probability the algorithm fails before $s \geq \operatorname{st.rank}(\boldsymbol{A})$ is at most

$$\sum\nolimits_{j=2}^{\infty} (7/8)^{8j} = \frac{(7/8)^{16}}{1 - (7/8)^8} < 0.2.$$

With constant probability, we obtain a set $\tau_\star$ with cardinality at least $\operatorname{st.rank}(\boldsymbol{A})/2$.

The cost of the procedure Norm-Reduce is dominated by the cost of the Pietsch factorization, which is $\widetilde{\mathrm{O}}(s^2 m + s^{3.5})$ for a fixed s. Summing over s and k, we find that the total cost of all the invocations of Norm-Reduce is dominated (up to logarithmic factors) by the cost of the final invocation, during which the parameter $s \leq 2\,|\tau_\star|$.

An estimate of the spectral norm of $\boldsymbol{A}_\tau$ can be obtained as a by-product of solving (3.2). Indeed, Proposition 2.3 and the discussion in Section 3.2 show that we can bound the spectral norm in terms of the parameter α and the objective value obtained in (3.2). □

5. The Bourgain–Tzafriri Theorem

Our proof of the Bourgain–Tzafriri theorem is almost identical in structure with the proof of the Kashin–Tzafriri theorem. This streamlined argument appears to be simpler than all previously published approaches, but it contains no significant conceptual innovations. Our discussion culminates in an algorithm remarkably similar to Algorithm 1.

5.1. **Preliminary results.** Suppose $\boldsymbol{A}$ is a standardized matrix with n columns. We will work instead with a related matrix $\boldsymbol{H} = \boldsymbol{A}^*\boldsymbol{A} - \mathbf{I}$, which is called the *hollow Gram matrix*. The advantage of considering the hollow Gram matrix is that we can perform column selection on $\boldsymbol{A}$ simply by reducing the norm of $\boldsymbol{H}$.

Proposition 5.1. *Suppose $\boldsymbol{A}$ is a standardized matrix with hollow Gram matrix $\boldsymbol{H}$. If τ is a set of column indices for which $\|\boldsymbol{H}_{\tau\times\tau}\| \le 0.5$, then $\kappa(\boldsymbol{A}_\tau) \le \sqrt{3}$.*

Proof. The hypothesis $\|\boldsymbol{H}_{\tau\times\tau}\| \le 0.5$ implies that the eigenvalues of $\boldsymbol{H}_{\tau\times\tau}$ lie in the range $[-0.5, 0.5]$. Since $\boldsymbol{H}_{\tau\times\tau} = \boldsymbol{A}_\tau^*\boldsymbol{A}_\tau - \mathbf{I}$, the eigenvalues of $\boldsymbol{A}_\tau^*\boldsymbol{A}_\tau$ fall in the interval $[0.5, 1.5]$. An equivalent condition is that $0.5 \le \|\boldsymbol{A}_\tau\boldsymbol{x}\|_2^2 \le 1.5$ whenever $\|\boldsymbol{x}\|_2 = 1$. We conclude that

$$\kappa(\boldsymbol{A}_\tau) = \max\left\{ \frac{\|\boldsymbol{A}_\tau\boldsymbol{x}\|_2}{\|\boldsymbol{A}_\tau\boldsymbol{y}\|_2} : \|\boldsymbol{x}\|_2 = \|\boldsymbol{y}\|_2 = 1 \right\} \le \sqrt{\frac{1.5}{0.5}} = \sqrt{3}.$$

Thus, a norm bound for $\boldsymbol{H}_{\tau\times\tau}$ yields a condition number bound for $\boldsymbol{A}_\tau$. □

As we mentioned before, random selection may reduce other norms even if it does not reduce the spectral norm. Define the natural norm on linear maps from ℓ_∞ to ℓ_1 by the formula

$$\|\boldsymbol{G}\|_{\infty\to 1} = \max\{\|\boldsymbol{G}\boldsymbol{x}\|_1 : \|\boldsymbol{x}\|_\infty = 1\}.$$

This norm is closely related to the cut norm, which plays a starring role in graph theory [AN04]. For a general $s \times s$ matrix $\boldsymbol{G}$, the best inequality between the $(\infty, 1)$ norm and the spectral norm is $\|\boldsymbol{G}\|_{\infty\to 1} \le s\,\|\boldsymbol{G}\|$. Rohn [Roh00] has established that there is a class of positive semidefinite, integer matrices for which it is NP-hard to determine the $(\infty, 1)$ norm within an absolute tolerance of $1/2$. Nevertheless, it can be approximated within a small relative factor in polynomial time [AN04].

The $(\infty, 1)$ norm decreases when we randomly sample a principal submatrix. The following result, which we establish in Appendix A.4, is a direct consequence of Rudelson and Vershynin's work on the cut norm of random submatrices [RV07, Thm. 1.5].

Theorem 5.2. *Suppose $\boldsymbol{A}$ is an n-column standardized matrix with hollow Gram matrix $\boldsymbol{H}$. Choose*

$$s \le \lceil \mathrm{c} \cdot \mathrm{st.rank}(\boldsymbol{A}) \rceil,$$

and draw a uniformly random subset σ with cardinality s from $\{1, 2, \dots, n\}$. Then

$$\mathbb{E}\,\|\boldsymbol{H}_{\sigma\times\sigma}\|_{\infty\to 1} \le \frac{s}{9}.$$

In particular, $\|\boldsymbol{H}_{\sigma\times\sigma}\|_{\infty\to 1} \le s/8$ with probability at least $1/9$.

To connect the $(\infty, 1)$ norm with the spectral norm, we call on the celebrated factorization of Grothendieck [Pis86, p. 56].

Theorem 5.3 (Grothendieck Factorization)**.** *Each matrix $\boldsymbol{G}$ can be factored as $\boldsymbol{G} = \boldsymbol{D}_1\boldsymbol{T}\boldsymbol{D}_2$ where*

(1) *$\boldsymbol{D}_i$ is a nonnegative, diagonal matrix with $\mathrm{trace}(\boldsymbol{D}_i^2) = 1$ for $i = 1, 2$, and*
(2) *$\|\boldsymbol{G}\|_{\infty\to 1} \le \|\boldsymbol{T}\| \le \mathrm{K_G}\,\|\boldsymbol{G}\|_{\infty\to 1}$.*

When $\boldsymbol{G}$ is Hermitian, we may take $\boldsymbol{D}_1 = \boldsymbol{D}_2$.

The precise value of the Grothendieck constant $\mathrm{K_G}$ remains an outstanding open question, but it is known to depend on the scalar field [Pis86, Sec. 5e].

- When the scalar field is real, $1.570 \le \pi/2 \le \mathrm{K_G}(\mathbb{R}) \le \pi/(2\log(1+\sqrt{2})) \le 1.783$.
- When the scalar field is complex, $1.338 \le \mathrm{K_G}(\mathbb{C}) \le 1.405$.

For positive semidefinite $\boldsymbol{G}$, the real (resp., complex) Grothendieck constant equals the square of the real (resp., complex) Pietsch constant because $\|\boldsymbol{B}^*\boldsymbol{B}\|_{\infty\to 1} = \|\boldsymbol{B}\|_{\infty\to 2}^2$.

The following proposition describes the role of the Grothendieck factorization in the selection of submatrices with controlled spectral norm.

Proposition 5.4. *Suppose $\boldsymbol{G}$ is an $s \times s$ Hermitian matrix. There is a set τ of column indices for which*

$$|\tau| \ge \frac{s}{2} \quad \textit{and} \quad \|\boldsymbol{G}_{\tau\times\tau}\| \le \frac{2\mathrm{K_G}}{s}\,\|\boldsymbol{G}\|_{\infty\to 1}\,.$$

Proof. Consider a Grothendieck factorization $\boldsymbol{G} = \boldsymbol{D}\boldsymbol{T}\boldsymbol{D}$, and identify $\tau = \{j : d_{jj}^2 \le s/2\}$. The remaining details echo the proof of Proposition 2.3. □

5.2. **Proof of Bourgain–Tzafriri.** Suppose $\boldsymbol{A}$ is a standardized matrix with n columns, and consider its hollow Gram matrix $\boldsymbol{H}$. Theorem 5.2 provides a set σ for which

$$|\sigma| \ge \mathrm{c}\cdot \mathrm{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \|\boldsymbol{H}_{\sigma\times\sigma}\|_{\infty\to 1} \le \frac{s}{8}.$$

Apply Proposition 5.4 to the $s \times s$ matrix $\boldsymbol{G} = \boldsymbol{H}_{\sigma\times\sigma}$ to obtain a further subset τ inside σ with

$$|\tau| \ge \frac{s}{2} \quad \text{and} \quad \|\boldsymbol{G}_{\tau\times\tau}\| \le \frac{2\mathrm{K_G}}{s}\,\|\boldsymbol{G}\|_{\infty\to 1}\,.$$

Since $2\mathrm{K_G} < 4$ and $\boldsymbol{H}_{\tau\times\tau} = \boldsymbol{G}_{\tau\times\tau}$, we determine that

$$|\tau| \ge \frac{\mathrm{c}}{2}\cdot \mathrm{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \|\boldsymbol{H}_{\tau\times\tau}\| \le 0.5.$$

In view of Proposition 5.1, we conclude $\kappa(\boldsymbol{A}_\tau) \le \sqrt{3}$.

Now, take another step back and notice that this here argument is nearly algorithmic. The random selection of σ can easily be implemented in practice, even though the proof does not specify the value of c. Given a Grothendieck factorization $\boldsymbol{G} = \boldsymbol{D}\boldsymbol{T}\boldsymbol{D}$, it is straightforward to identify the subset τ. The challenge, as before, is to produce the factorization.

5.3. **Grothendieck factorization via convex optimization.** As with the Pietsch factorization, the Grothendieck factorization can be identified from the solution to a convex program.

Theorem 5.5. *Suppose $\boldsymbol{G}$ is Hermitian. The factorization $\boldsymbol{G} = \boldsymbol{D}\boldsymbol{T}\boldsymbol{D}$ satisfies $\|\boldsymbol{T}\| \le \alpha$ if and only if $\boldsymbol{D}$ satisfies*

$$\lambda_{\max}\begin{bmatrix} -\alpha\boldsymbol{D}^2 & \boldsymbol{G} \\ \boldsymbol{G} & -\alpha\boldsymbol{D}^2 \end{bmatrix} \le 0. \tag{5.1}$$

In particular, if no $\boldsymbol{D}$ verifies this bound, then no factorization $\boldsymbol{G} = \boldsymbol{D}\boldsymbol{T}\boldsymbol{D}$ admits $\|\boldsymbol{T}\| \le \alpha$.

Proof. To check the forward implication, we essentially repeat the argument we used in Theorem 3.1 for the Pietsch case. This reasoning yields the pair of relations

$$\boldsymbol{G} - \alpha\boldsymbol{D}^2 \preccurlyeq \boldsymbol{0} \quad \text{and} \quad -\boldsymbol{G} - \alpha\boldsymbol{D}^2 \preccurlyeq \boldsymbol{0}.$$

Together, these two relations are equivalent with (5.1) because

$$\begin{bmatrix} -\alpha\boldsymbol{D}^2 & \boldsymbol{G} \\ \boldsymbol{G} & -\alpha\boldsymbol{D}^2 \end{bmatrix} = \frac{1}{2}\begin{bmatrix} \mathbf{I} & \mathbf{I} \\ -\mathbf{I} & \mathbf{I} \end{bmatrix}^* \begin{bmatrix} \boldsymbol{G} - \alpha\boldsymbol{D}^2 & \\ & -\boldsymbol{G} - \alpha\boldsymbol{D}^2 \end{bmatrix} \begin{bmatrix} \mathbf{I} & \mathbf{I} \\ -\mathbf{I} & \mathbf{I} \end{bmatrix}.$$

To prove the reverse implication, we assume that (5.1) holds. First, we must check that $d_{jj} = 0$ implies that $\boldsymbol{g}_j = \boldsymbol{0}$. To verify this claim, observe that

$$0 \ge \begin{bmatrix} \alpha \\ \boldsymbol{g}_j \end{bmatrix}^* \begin{bmatrix} 0 & \boldsymbol{g}_j^* \\ \boldsymbol{g}_j & -\alpha\boldsymbol{D}^2 \end{bmatrix} \begin{bmatrix} \alpha \\ \boldsymbol{g}_j \end{bmatrix} = \alpha\left(2\,\|\boldsymbol{g}_j\|_2^2 - \boldsymbol{g}_j^*\boldsymbol{D}^2\boldsymbol{g}_j\right) \ge \alpha\,\|\boldsymbol{g}_j\|_2^2$$

because $\operatorname{trace}(\boldsymbol{D}^2) = 1$. Therefore, we may construct a Grothendieck factorization $\boldsymbol{G} = \boldsymbol{DTD}$ with $\|\boldsymbol{T}\| \leq \alpha$ by setting $\boldsymbol{T} = \boldsymbol{D}^\dagger \boldsymbol{G} \boldsymbol{D}^\dagger$. □

This discussion leads us to frame the eigenvalue minimization problem

$$\min \ \lambda_{\max} \begin{bmatrix} -\alpha \boldsymbol{F} & \boldsymbol{G} \\ \boldsymbol{G} & -\alpha \boldsymbol{F} \end{bmatrix} \qquad \text{subject to} \qquad \operatorname{trace}(\boldsymbol{F}) = 1, \ \boldsymbol{F} \text{ diagonal}, \ \boldsymbol{F} \geq \boldsymbol{0} \tag{5.2}$$

Owing to Theorem 5.5, there is a factorization $\boldsymbol{G} = \boldsymbol{DTD}$ with $\|\boldsymbol{T}\| \leq \alpha$ if and only if the value of (5.2) is nonpositive.

As in Section 3.2, we can easily construct Grothendieck factorizations from (imprecise) solutions to the problem (5.2). The proof of Bourgain–Tzafriri suggests that an appropriate value for the parameter $\alpha = s/4$. Furthermore, we do not need to solve (5.2) to optimality to obtain the required information. Indeed, it suffices to produce a feasible point with an objective value of $\mathrm{O}(1)$.

To solve (5.2) in practice, we again propose the Entropic Mirror Descent algorithm [BT03]. Appendix B describes the application to this problem. To provide a concrete bound on the computational cost, we remark that, when $\boldsymbol{A}_\tau$ has dimension $m \times s$, forming $\boldsymbol{G} = \boldsymbol{A}_\tau^* \boldsymbol{A}_\tau - \mathbf{I}$ costs at most $\mathrm{O}(s^2 m)$, and Alizadeh's interior-point method [Ali95] requires $\widetilde{\mathrm{O}}(s^{3.5})$ time.

Remark 5.6. *For symmetric $\boldsymbol{G}$, Theorem 5.3 shows that the norm $\|\boldsymbol{G}\|_{\infty \to 1}$ is approximated within a factor $\mathrm{K_G}$ by the least α for which* (5.2) *has a nonpositive value. A natural reformulation of* (5.2) *can identify this value of α automatically (cf. Section 3.3). For nonsymmetric $\boldsymbol{G}$, similar optimization problems arise. These ideas yield new approximation algorithms for the $(\infty, 1)$ norm.*

5.4. **An algorithm for Bourgain–Tzafriri.** We are prepared to state our algorithm for producing the set τ described by the Bourgain–Tzafriri theorem. The procedure appears as Algorithm 2 on page 11. Note the striking similarity with Algorithm 1. The following result describes the performance of the algorithm. We omit the proof, which parallels that of Theorem 4.1.

Theorem 5.7. *Suppose $\boldsymbol{A}$ is an $m \times n$ standardized matrix. With probability at least $3/4$, Algorithm 2 produces a set $\tau = \tau_\star$ of column indices for which*

$$|\tau| \geq \mathrm{c} \cdot \operatorname{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \kappa(\boldsymbol{A}_\tau) \leq \sqrt{3}.$$

The computational cost is bounded by $\widetilde{\mathrm{O}}(|\tau|^2 \, m + |\tau|^{3.5})$.

6. Future Directions

After the initial work [BT87], additional research has clarified the role of the stable rank. We highlight a positive result of Vershynin [Ver01, Cor. 7.1] and a negative result of Szarek [Sza90, Thm. 1.2] which together imply that the stable rank describes *precisely* how large a well-conditioned column submatrix can in general exist. See [Ver01, Sec. 5] for a more detailed discussion.

Theorem 6.1 (Vershynin 2001)**.** *Fix $\varepsilon > 0$. For each matrix $\boldsymbol{A}$, there is a set τ of column indices for which*

$$|\tau| \geq (1 - \varepsilon) \cdot \operatorname{st.rank}(\boldsymbol{A}) \quad \text{and} \quad \kappa(\boldsymbol{A}_\tau) \leq \mathrm{C}(\varepsilon).$$

Theorem 6.2 (Szarek)**.** *There is a sequence $\{\boldsymbol{A}(n)\}$ of matrices of increasing dimension for which*

$$|\tau| = \operatorname{st.rank}(\boldsymbol{A}) \quad \Longrightarrow \quad \kappa(\boldsymbol{A}_\tau) = \omega(1).$$

Vershynin's proof constructs the set τ in Theorem 6.1 with a complicated iteration that interleaves the Kashin–Tzafriri theorem and the Bourgain–Tzafriri theorem. We believe that the argument can be simplified substantially and developed into a column selection algorithm. This achievement might lead to a new method for performing rank-revealing factorizations, which could have a significant impact on the practice of numerical linear algebra.

Algorithm 1: Constructive version of Kashin–Tzafriri Theorem

KT($\boldsymbol{A}$)
Input: Standardized matrix $\boldsymbol{A}$ with n columns
Output: A subset $\tau_\star$ of $\{1, 2, \dots, n\}$
Description: Produces $\tau_\star$ such that $|\tau_\star| \geq \text{st.rank}(\boldsymbol{A})/2$ and $\|\boldsymbol{A}_\tau\| \leq 15$ w.p. 4/5

$\tau_\star = \{1\}$
for $s = 4, 8, 16, \dots, n$
  for $k = 1, 2, 3, \dots, 8 \log_2 s$
    $\tau =$ NORM-REDUCE($\boldsymbol{A}, s$)
    if $\|\boldsymbol{A}_\tau\| \leq 15$ **then** $\tau_\star = \tau$ and **break**
  if $|\tau_\star| < s$ **then exit**

NORM-REDUCE($\boldsymbol{A}$, s)
Input: Standardized matrix $\boldsymbol{A}$ with n columns, a parameter s
Output: A subset τ of $\{1, 2, \dots, n\}$

Draw a uniformly random set σ with cardinality s from $\{1, 2, \dots, n\}$
Solve (3.2) with $\boldsymbol{B} = \boldsymbol{A}_\sigma$ and $\alpha = 8\text{K}_\text{P}\sqrt{s}$ to obtain a factorization $\boldsymbol{B} = \boldsymbol{T}\boldsymbol{D}$
Return $\tau = \{j \in \sigma : d_{jj}^2 \leq 2/s\}$

Algorithm 2: Constructive version of Bourgain–Tzafriri Theorem

BT($\boldsymbol{A}$)
Input: Standardized matrix $\boldsymbol{A}$ with n columns
Output: A subset $\tau_\star$ of $\{1, 2, \dots, n\}$
Description: Produces $\tau_\star$ such that $|\tau_\star| \geq \text{st.rank}(\boldsymbol{A})/2$ and $\kappa(\boldsymbol{A}_\tau) \leq \sqrt{3}$ w.p. 3/4

$\tau_\star = \{1\}$
for $s = 4, 8, 16, \dots, n$
  for $k = 1, 2, 3, \dots, 8 \log_2 s$
    $\tau =$ COND-REDUCE($\boldsymbol{A}, s$)
    if $\kappa(\boldsymbol{A}_\tau) \leq \sqrt{3}$ **then** $\tau_\star = \tau$ and **break**
  if $|\tau_\star| < s$ **then exit**

COND-REDUCE($\boldsymbol{A}$, s)
Input: Standardized matrix $\boldsymbol{A}$ with n columns, a parameter s
Output: A subset τ of $\{1, 2, \dots, n\}$

Draw a uniformly random set σ with cardinality s from $\{1, 2, \dots, n\}$
Solve (5.2) with $\boldsymbol{G} = \boldsymbol{A}_\sigma^* \boldsymbol{A}_\sigma - \mathbf{I}$ and $\alpha = s/4$ to obtain factorization $\boldsymbol{G} = \boldsymbol{D}\boldsymbol{T}\boldsymbol{D}$
Return $\tau = \{j \in \sigma : d_{jj}^2 \leq 2/s\}$

Appendix A. Random Reduction of Norms

How does the norm of a matrix change when we pass to a random submatrix? This question has great importance in modern functional analysis, but it also has implications for the design of algorithms. This appendix describes some general results on how random selection reduces the $(\infty, 2)$ norm and the $(\infty, 1)$ norm. We also specialize these results to the structured matrices that appear in the proofs of Theorem 1.1 and Theorem 1.2.

A.1. Random Coordinate Models.

We begin with two standard models for selecting random submatrices, and we describe how these models are related for an important class of matrix norms.

A matrix norm is *monotonic* if the norm of a matrix exceeds the norm of every (rectangular) submatrix. More precisely, the norm $\|\!|\!|\cdot|\!|\!|$ is monotonic if

$$\|\!|\!| \boldsymbol{P}\boldsymbol{A}\boldsymbol{P}' |\!|\!| \leq \|\!|\!| \boldsymbol{A} |\!|\!|$$

for each matrix $\boldsymbol{A}$ and each pair $\boldsymbol{P}, \boldsymbol{P}'$ of diagonal (i.e., coordinate) projectors. The basic example of a monotonic matrix norm is the natural norm on operators from ℓ_p to ℓ_q with p, q in $[1, \infty]$, which is defined as

$$\|\boldsymbol{A}\|_{p \to q} = \max\{\|\boldsymbol{A}\boldsymbol{x}\|_q : \|\boldsymbol{x}\|_p = 1\}.$$

Fix a number δ in $[0, 1]$, and denote by $\boldsymbol{P}_\delta$ a random $n \times n$ diagonal matrix where exactly $s = \lfloor \delta n \rfloor$ entries equal one and the rest equal zero. This matrix can be viewed as a projector onto a random set of s coordinates. Therefore, we may treat $\boldsymbol{A}\boldsymbol{P}_\delta$ as a random s-column submatrix of $\boldsymbol{A}$ by ignoring the zeroed columns. Although this model is conceptually appealing, it can be difficult to analyze because of the dependencies among coordinates.

Let us introduce a simpler model for selecting random coordinates. We denote by $\boldsymbol{R}_\delta$ a random $n \times n$ diagonal matrix whose entries are independent 0–1 random variables with common mean δ. This matrix is a projector onto a random set of coordinates with *average* cardinality δn.

There is a basic result connecting these two models. The statement here follows directly from the argument in [Tro08, Lem. 14].

Proposition A.1 (Poissonization)**.** *Let $\|\!|\!|\cdot|\!|\!|$ be a monotonic matrix norm. For each matrix $\boldsymbol{A}$ with n columns, it holds that*

$$\mathbb{E}\, \|\!|\!| \boldsymbol{A}\boldsymbol{P}_\delta |\!|\!| \leq 2\, \mathbb{E}\, \|\!|\!| \boldsymbol{A}\boldsymbol{R}_\delta |\!|\!| .$$

For each $n \times n$ matrix $\boldsymbol{H}$, it holds that

$$\mathbb{E}\, \|\!|\!| \boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta |\!|\!| \leq 2\, \mathbb{E}\, \|\!|\!| \boldsymbol{R}_\delta \boldsymbol{H} \boldsymbol{R}_\delta |\!|\!| .$$

A.2. Reduction of the $(\infty, 2)$ norm.

We begin with a general result on the $(\infty, 2)$ norm of a uniformly random set of columns drawn from a fixed matrix. The basic argument appears already in the work of Bourgain and Tzafriri [BT91, Thm. 1.1], but modern proofs are a little simpler. (See [Ver06, Lem. 2.3], for example.) The version here offers especially good constants.

Theorem A.2. *Fix $\delta \in [0, 1]$, and suppose $\boldsymbol{A}$ is a matrix with n columns. Then*

$$\mathbb{E}\, \|\boldsymbol{A}\boldsymbol{R}_\delta\|_{\infty \to 2} \leq \sqrt{2\delta(1-\delta)}\, \|\boldsymbol{A}\|_{\mathrm{F}} + \delta\, \|\boldsymbol{A}\|_{\infty \to 2} .$$

We postpone the argument to the next section so we may note a corollary that appears as a key step in the proof of the Kashin–Tzafriri theorem.

Corollary A.3. *Suppose $\boldsymbol{A}$ is a standardized matrix with n columns. Choose $s \leq \lceil 2 \operatorname{st.rank}(\boldsymbol{A}) \rceil$, and write $\delta = s/n$. Then*

$$\mathbb{E}\, \|\boldsymbol{A}\boldsymbol{P}_\delta\|_{\infty \to 2} \leq 7\sqrt{s}.$$

Proof. Owing to the standardization, $1 \le \operatorname{st.rank}(\boldsymbol{A}) = n/\|\boldsymbol{A}\|^2$. It follows that

$$\delta \le \frac{2\operatorname{st.rank}(\boldsymbol{A}) + 1}{n} \le \frac{3\operatorname{st.rank}(\boldsymbol{A})}{n} = \frac{3}{\|\boldsymbol{A}\|^2}.$$

Apply the Poissonization result, Proposition A.1, to see that

$$\mathbb{E}\,\|\boldsymbol{A}\boldsymbol{P}_\delta\|_{\infty\to 2} \le 2\,\mathbb{E}\,\|\boldsymbol{A}\boldsymbol{R}_\delta\|_{\infty\to 2}\,.$$

Theorem A.2 yields

$$\mathbb{E}\,\|\boldsymbol{A}\boldsymbol{P}_\delta\|_{\infty\to 2} \le 2\sqrt{2\delta}\,\|\boldsymbol{A}\|_{\mathrm{F}} + 2\delta\,\|\boldsymbol{A}\|_{\infty\to 2}\,.$$

Since $\boldsymbol{A}$ has n unit-norm columns, it holds that $\|\boldsymbol{A}\|_{\mathrm{F}} = \sqrt{n}$. We also have the general bound $\|\boldsymbol{A}\|_{\infty\to 2} \le \sqrt{n}\,\|\boldsymbol{A}\|$. Therefore,

$$\mathbb{E}\,\|\boldsymbol{A}\boldsymbol{P}_\delta\|_{\infty\to 2} \le 2\sqrt{2\delta n} + 2\delta\sqrt{n}\,\|\boldsymbol{A}\| = 2\sqrt{s}\left[\sqrt{2} + \sqrt{\delta}\,\|\boldsymbol{A}\|\right].$$

Introduce the bound on δ and make a numerical estimate to complete the proof. □

A.3. **Proof of Theorem A.2.** We must bound the quantity

$$E = \mathbb{E}\,\|\boldsymbol{A}\boldsymbol{R}_\delta\|_{\infty\to 2}\,.$$

It turns out that it is easier to work with the $(2,1)$ norm, which is dual to the $(\infty, 2)$ norm, because there are some special methods that apply. Rewrite the expression as

$$E = \mathbb{E}\,\|\boldsymbol{R}_\delta\boldsymbol{A}^*\|_{2\to 1} = \mathbb{E}\max_{\|\boldsymbol{x}\|_2=1}\sum\nolimits_{j=1}^{n}\delta_j\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle|$$

where $\{\delta_j\}$ is a sequence of independent 0–1 random variables with common mean δ. In the sequel, we simplify notation by omitting the restriction on the vector $\boldsymbol{x}$ and the limits from the sum.

The next step is to center and symmetrize the selectors. First, add and subtract the mean of each term from the sum and use the subadditivity of the maximum to obtain

$$\begin{aligned} E &\le \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j(\delta_j - \delta)\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle| + \max_{\boldsymbol{x}}\sum\nolimits_j \delta\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle| \\ &= \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j(\delta_j - \delta)\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle| + \delta\,\|\boldsymbol{A}^*\|_{2\to 1} \\ &= \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j(\delta_j - \delta)\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle| + \delta\,\|\boldsymbol{A}\|_{\infty\to 2}\,. \end{aligned}$$

We focus on the first term, which we abbreviate by the letter F. Let $\{\delta_j'\}$ be an independent copy of the sequence $\{\delta_j\}$. Jensen's inequality allows that

$$\begin{aligned} F &= \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j(\delta_j - \mathbb{E}\,\delta_j')\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle| \\ &\le \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j(\delta_j - \delta_j')\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle|\,. \end{aligned}$$

Observe that $\{\delta_j - \delta_j'\}$ is a sequence of independent, symmetric random variables. Thus, we may multiply each one by a random sign without changing the expectation [LT91, Lem. 6.3]. That is,

$$F \le \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j \varepsilon_j(\delta_j - \delta_j')\,|\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle|$$

where $\{\varepsilon_j\}$ is a sequence of independent Rademacher (i.e., uniform ± 1) random variables.

Now, we invoke a specific type of Rademacher comparison [LT91, Thm. 4.12 et seq.] to remove the absolute values from the inner product:

$$F \le \mathbb{E}\max_{\boldsymbol{x}}\sum\nolimits_j \varepsilon_j(\delta_j - \delta_j')\,\langle \boldsymbol{a}_j,\ \boldsymbol{x}\rangle = \mathbb{E}\max_{\boldsymbol{x}}\left\langle\sum\nolimits_j \varepsilon_j(\delta_j - \delta_j')\boldsymbol{a}_j,\ \boldsymbol{x}\right\rangle.$$

Since $\boldsymbol{x}$ ranges over the ℓ_2 unit sphere, we reach

$$F \le \mathbb{E}\left\|\sum\nolimits_j \varepsilon_j(\delta_j - \delta_j')\boldsymbol{a}_j\right\|_2.$$

The remaining expectations are elementary. First, apply Hölder's inequality to obtain

$$F \le \left(\mathbb{E} \left\| \sum\nolimits_j \varepsilon_j (\delta_j - \delta_j') \boldsymbol{a}_j \right\|_2^2 \right)^{1/2}.$$

Compute the expectation with respect to $\{\varepsilon_j\}$ and then with respect to $\{\delta_j\}$ and $\{\delta_j'\}$.

$$F \le \left(\mathbb{E} \sum\nolimits_j (\delta_j - \delta_j')^2 \left\| \boldsymbol{a}_j \right\|_2^2 \right)^{1/2} = \left(\sum\nolimits_j 2\delta(1-\delta) \left\| \boldsymbol{a}_j \right\|_2^2 \right)^{1/2} = \sqrt{2\delta(1-\delta)} \left\| \boldsymbol{A} \right\|_{\mathrm{F}}.$$

Introduce this bound on F into the bound on E to conclude that

$$\mathbb{E} \left\| \boldsymbol{A} \boldsymbol{P}_\delta \right\|_{\infty \to 2} \le \sqrt{2\delta(1-\delta)} \left\| \boldsymbol{A} \right\|_{\mathrm{F}} + \delta \left\| \boldsymbol{A} \right\|_{\infty \to 2}.$$

This is the advertised estimate.

A.4. **Reduction of the $(\infty, 1)$ norm.** The impact of random selection on the $(\infty, 1)$ norm has already received some attention in the theoretical computer science literature because of a connection with graph cuts. The following result of Rudelson and Vershynin contains detailed information on the $(\infty, 1)$ norm of a random principal submatrix. The statement involves an auxiliary norm

$$\left\| \boldsymbol{H} \right\|_{\mathrm{col}} = \sum\nolimits_j \left\| \boldsymbol{H} \mathbf{e}_j \right\|_2,$$

where $\{\mathbf{e}_j\}$ is the set of standard basis vectors. In words, we sum the Euclidean norms of the columns of the matrix.

Theorem A.4 (Rudelson–Vershynin)**.** *Fix $\delta \in [0, 1]$, and suppose $\boldsymbol{H}$ is an $n \times n$ matrix. Then*

$$\mathbb{E} \left\| \boldsymbol{R}_\delta \boldsymbol{H} \boldsymbol{R}_\delta \right\|_{\infty \to 1} \le \mathrm{C} \left[\delta^2 \left\| \boldsymbol{H} - \operatorname{diag}(\boldsymbol{H}) \right\|_{\infty \to 1} + \delta^{3/2} \left(\left\| \boldsymbol{H} \right\|_{\mathrm{col}} + \left\| \boldsymbol{H}^* \right\|_{\mathrm{col}} \right) + \delta \left\| \operatorname{diag}(\boldsymbol{H}) \right\|_{\infty \to 1} \right].$$

Theorem A.4 is established with the same methods as Theorem A.2, along with an additional decoupling argument [BT91, Prop. 1.9]. We rely on the following corollary in our proof of the Bourgain–Tzafriri theorem.

Corollary A.5. *Suppose $\boldsymbol{A}$ is an n-column standardized matrix with hollow Gram matrix $\boldsymbol{H} = \boldsymbol{A}^* \boldsymbol{A} - \mathbf{I}$. Choose $s \le \lceil \mathrm{c} \cdot \mathrm{st.rank}(\boldsymbol{A}) \rceil$, and write $\delta = s/n$. Then*

$$\mathbb{E} \left\| \boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta \right\|_{\infty \to 1} \le \frac{s}{9}.$$

Proof. Suppose $\boldsymbol{A}$ is a standardized matrix with n columns, and define its $n \times n$ hollow Gram matrix $\boldsymbol{H}$. Observe that the $(\infty, 1)$ norm of $\boldsymbol{H}$ satisfies the bound

$$\left\| \boldsymbol{H} \right\|_{\infty \to 1} \le n \left\| \boldsymbol{H} \right\| \le n \max\{ \left\| \boldsymbol{A}^* \boldsymbol{A} \right\| - 1, 1 \} \le n \left\| \boldsymbol{A} \right\|^2.$$

Meanwhile, the $\left\| \cdot \right\|_{\mathrm{col}}$ norm satisfies

$$\left\| \boldsymbol{H} \right\|_{\mathrm{col}} < \left\| \boldsymbol{A}^* \boldsymbol{A} \right\|_{\mathrm{col}} = \sum\nolimits_j \left\| \boldsymbol{A}^* \boldsymbol{a}_j \right\|_2 \le n \left\| \boldsymbol{A} \right\|.$$

These facts play a central role in the calculation.

To continue, invoke the Poissonization result, Proposition A.1, which yields

$$\mathbb{E} \left\| \boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta \right\|_{\infty \to 1} \le 2 \left\| \boldsymbol{R}_\delta \boldsymbol{H} \boldsymbol{R}_\delta \right\|_{\infty \to 1}.$$

Theorem A.4 provides that

$$\mathbb{E} \left\| \boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta \right\|_{\infty \to 1} \le \mathrm{C} \left[\delta^2 \left\| \boldsymbol{H} \right\|_{\infty \to 1} + \delta^{3/2} \left\| \boldsymbol{H} \right\|_{\mathrm{col}} \right]$$

where we have applied the facts that $\boldsymbol{H}$ is Hermitian and has a zero diagonal. The two norm bounds result in additional simplifications:

$$\left\| \boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta \right\|_{\infty \to 1} \le \mathrm{C} \left[\delta^2 n \left\| \boldsymbol{A} \right\|^2 + \delta^{3/2} n \left\| \boldsymbol{A} \right\| \right] = \mathrm{C} s \left[\delta \left\| \boldsymbol{A} \right\|^2 + \delta^{1/2} \left\| \boldsymbol{A} \right\| \right].$$

Since $\boldsymbol{A}$ has unit-norm columns, $\operatorname{st.rank}(\boldsymbol{A}) = n/\|\boldsymbol{A}\|^2$. As a result, $\delta = s/n \leq \mathrm{c}/\|\boldsymbol{A}\|^2$. By fixing a sufficiently small constant c, we can ensure that

$$\|\boldsymbol{P}_\delta \boldsymbol{H} \boldsymbol{P}_\delta\|_{\infty\to 1} \leq \frac{s}{9},$$

the advertised bound. □

Appendix B. Entropic Mirror Descent

The algorithms for the Kashin–Tzafriri theorem and the Bourgain–Tzafriri theorem both require the solution to a convex minimization problem over the probability simplex. It is important to have a practical algorithm for approaching these optimizations. To that end, we briefly describe a simple, elegant method called Entropic Mirror Descent [BT03]. We then explain how to apply this technique to the specific objective functions that arise in our work.

B.1. **Convex analysis.** Let $\mathbb{E}$ be a Euclidean space, i.e., a vector space equipped with a real-linear inner product. Let Ω be a convex subset of $\mathbb{E}$, and consider a convex function $J : \Omega \to \mathbb{R}$. The *subdifferential* $\partial J(\boldsymbol{f})$ contains each vector $\boldsymbol{\theta} \in \mathbb{E}^*$ that satisfies the inequalities

$$J(\boldsymbol{h}) - J(\boldsymbol{f}) \geq \langle \boldsymbol{\theta},\ \boldsymbol{h} - \boldsymbol{f} \rangle \quad \text{for all } \boldsymbol{h} \in \Omega.$$

The elements of the subdifferential are called *subgradients*. They describe the directions and rates of ascent of the function J at the point $\boldsymbol{f}$. When J is differentiable at $\boldsymbol{f}$, the gradient is the unique subgradient.

The *Lipschitz constant* of the function J with respect to a norm $\|\!|\cdot|\!\|$ is defined to be the least number L for which

$$|J(\boldsymbol{h}) - J(\boldsymbol{f})| \leq L\,|\!|\!|\boldsymbol{h} - \boldsymbol{f}|\!|\!| \quad \text{for all } \boldsymbol{h}, \boldsymbol{f} \in \Omega.$$

It can be shown [Roc70, Thm. 24.7] that

$$L = \sup\{ |\!|\!|\boldsymbol{\theta}|\!|\!|_* : \boldsymbol{\theta} \in \partial J(\boldsymbol{f}),\ \boldsymbol{f} \in \Omega \}.$$

where $\|\!|\cdot|\!\|_*$ is the dual norm.

B.2. **Interior subgradient methods.** Consider the (nonsmooth) convex program

$$\min\ J(\boldsymbol{f}) \qquad \text{subject to} \qquad \boldsymbol{f} \in \Omega.$$

Subgradient information can be used to solve this problem, but caution is necessary because the negative subgradient is not necessarily a direction of descent. As a result, subgradient methods are typically *nonmonotone*, which means that the value of the objective function can (and often will) increase. It is also common for subgradient methods to produce iterates outside the constraint set. The classical remedy is to project each iterate back onto the constraint set. This idea succeeds, but it leads to zigzagging phenomena.

Interior subgradient methods [BT03] are designed to eliminate some of the problematic behavior that projected subgradient methods exhibit. To develop an interior subgradient method, we need a divergence measure that is tailored to the constraint set. At each iteration, we perform two steps:

(1) At the current iterate $\boldsymbol{f}$, compute a subgradient $\boldsymbol{\theta} \in \partial J(\boldsymbol{f})$ to linearize the objective function:
$$J(\boldsymbol{h}) \approx J(\boldsymbol{f}) + \langle \boldsymbol{\theta},\ \boldsymbol{h} - \boldsymbol{f} \rangle$$
(2) Penalize the linearization with the divergence $D(\cdot; \boldsymbol{f})$ from the current iterate, scaled by a (large) parameter β^{-1}. Minimize this auxiliary function to produce a new iterate $\boldsymbol{f}'$:
$$\boldsymbol{f}' \in \operatorname*{arg\,min}_{\boldsymbol{h} \in \Omega} \left\{ J(\boldsymbol{f}) + \langle \boldsymbol{\theta},\ \boldsymbol{h} - \boldsymbol{f} \rangle + \beta^{-1} D(\boldsymbol{h}; \boldsymbol{f}) \right\}.$$

Algorithm 3: Entropic Mirror Descent

EMD(J, s, T)
Input: Objective function J, dimension s, number T of iterations
Output: Approximate minimizer $\boldsymbol{f}$ of J

$\boldsymbol{f}^{(1)} = s^{-1}\mathbf{e}$ { Initialize with uniform density }
for $t = 1$ **to** T
  Find $\boldsymbol{\theta} \in \partial J(\boldsymbol{f}^{(t)})$ { Compute subgradient }
  $\beta = \sqrt{\frac{2 \log s}{T \|\boldsymbol{\theta}\|_\infty^2}}$ { Compute step size }
  $\boldsymbol{h} = \boldsymbol{f}^{(t)} \cdot \exp(-\beta \boldsymbol{\theta})$ { Reweight current iterate }
  $\boldsymbol{f}^{(t+1)} = \boldsymbol{h} / \operatorname{trace}(\boldsymbol{h})$ { Rescale to obtain next iterate }
end *for*
Return $\boldsymbol{f} \in \arg\min_t J(\boldsymbol{f}^{(t)})$

The divergence penalty serves two purposes. First, it ensures that the next iterate is close to the previous iterate, which is essential because the linearization is only useful locally. Second, it simultaneously prevents the iterates from getting too close to the boundary of the constraint set. With a careful choice of the parameter β, we can guarantee progress toward the optimum set, at least on average.

B.3. **Optimization on the probability simplex.** The Entropic Mirror Descent (EMD) algorithm of Beck and Teboulle [BT03] is a specific instance of the interior subgradient method that is designed for minimizing convex functions over the probability simplex, the set defined by

$$\Delta_s = \{\boldsymbol{f} \in \mathbb{R}^s : \operatorname{trace}(\boldsymbol{f}) = 1,\ \boldsymbol{f} \geq \mathbf{0}\}.$$

A natural divergence measure for this set is the relative entropy function:

$$D(\boldsymbol{h}; \boldsymbol{f}) = \sum\nolimits_{j=1}^{s} h_j \log\left(\frac{h_j}{f_j}\right).$$

An amazing feature of the resulting interior subgradient method is that the optimization in the second step has a closed form:

$$h_j = \frac{f_j \exp(-\beta\theta_j)}{\sum_j f_j \exp(-\beta\theta_j)} \quad \text{for } j = 1, 2, \ldots, s.$$

Algorithm 3 describes the procedure that arises from these choices. Beck and Teboulle have established an elegant efficiency estimate [BT03, Thm. 4.2] for this method.

Theorem B.1 (Efficiency of EMD)**.** *Let $J : \Delta_s \to \mathbb{R}$ be a convex function whose Lipschitz constant with respect to the ℓ_1 norm is L. The approximate minimizer $\boldsymbol{f}$ generated by Algorithm 3 satisfies*

$$J(\boldsymbol{f}) - J(\boldsymbol{f}_\star) \leq \sqrt{\frac{2L^2 \log s}{T}}$$

where $\boldsymbol{f}_\star$ is a minimizer of J.

Algorithm 3 succeeds with a wide range of step sizes. In particular, when the total number T of iterations is unknown, we may compute the step size using the current iteration number t:

$$\beta = \sqrt{\frac{2 \log s}{t \, \|\boldsymbol{\theta}\|_\infty^2}}.$$

This choice increases the right-hand side of the efficiency estimate by a logarithmic factor.

B.4. **Pietsch factorization via EMD.** Suppose $\boldsymbol{B}$ is a matrix with s columns. We can rephrase the Pietsch factorization problem (3.2) as an optimization over the probability simplex. Define the linear operator

$$\operatorname{diag} : \mathbb{R}^s \to \mathbb{R}^{s \times s}$$

that maps vectors to diagonal matrices in the obvious way. We can write the convex program as

$$\min \ \lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2 \operatorname{diag}(\boldsymbol{f})) \qquad \text{subject to} \qquad \boldsymbol{f} \in \Delta_s. \tag{B.1}$$

Abbreviate the objective function $J : \Delta_s \to \mathbb{R}$. We can evidently apply EMD to complete the optimization once we find a way to compute subgradients.

We use methods from the convex analysis of Hermitian matrices to determine the subdifferential of the objective function [Lew96]. Let $\boldsymbol{A}$ be an Hermitian matrix. Then

$$\partial\lambda_{\max}(\boldsymbol{A}) = \operatorname{conv}\{\boldsymbol{u}\boldsymbol{u}^* : \boldsymbol{A}\boldsymbol{u} = \lambda_{\max}(\boldsymbol{A})\boldsymbol{u},\ \|\boldsymbol{u}\|_2 = 1\}.$$

In words, the subdifferential of the maximum eigenvalue function at $\boldsymbol{A}$ is the convex hull of all rank-one projectors whose range lies in the top eigenspace of $\boldsymbol{A}$. According to [Roc70, Thm. 23.9], we have

$$\partial J(\boldsymbol{f}) = \{-\alpha^2 \operatorname{diag}^*(\boldsymbol{\Theta}) : \boldsymbol{\Theta} \in \partial\lambda_{\max}(\boldsymbol{B}^*\boldsymbol{B} - \alpha^2 \operatorname{diag}(\boldsymbol{f}))\}.$$

where the adjoint map $\operatorname{diag}^* : \mathbb{R}^{s\times s} \to \mathbb{R}^s$ extracts the diagonal of a matrix. In particular, we may construct a subgradient $\boldsymbol{\theta} \in \partial J(\boldsymbol{f})$ from a normalized maximal eigenvector $\boldsymbol{u}$ of the matrix $\boldsymbol{B}^*\boldsymbol{B} - \alpha^2 \operatorname{diag}(\boldsymbol{f})$ using the formula

$$\boldsymbol{\theta} = -\alpha^2 \operatorname{diag}^*(\boldsymbol{u}\boldsymbol{u}^*) = -\alpha^2 |\boldsymbol{u}|^2$$

where $|\cdot|^2$ denotes the componentwise squared magnitude of a vector.

In summary, we can evaluate the objective function $J(\boldsymbol{f})$ and simultaneously obtain a subgradient $\boldsymbol{\theta} \in \partial J(\boldsymbol{f})$ from an eigenvector calculation plus some lower-order operations. Note that the standard methods for producing a single eigenvector, such as the Lanczos algorithm and its variants [GVL96, Ch. 9], require access to the matrix only through its action on vectors. It is therefore preferable in some settings—for example, when $\boldsymbol{B}$ is sparse—not to form the matrix $\boldsymbol{B}^*\boldsymbol{B}$.

Eigenvector computation is a primitive in every numerical linear package, so it is reasonable to assume that high-precision eigenvectors are available. In any case, slight variants of EMD will work with approximate subgradients, provided they are computed to sufficient precision. A simple analysis supporting this claim does not seem to be available in the optimization literature, but see [Kal07, Ch. 6] for related work.

We can bound the Lipschitz constant of J with respect to the ℓ_1 norm just by considering subgradients of the form $\boldsymbol{\theta} = -\alpha^2|\boldsymbol{u}|^2$ because their convex hull yields the complete subdifferential. Since the eigenvector $\boldsymbol{u}$ is normalized,

$$\|\boldsymbol{\theta}\|_\infty = \alpha^2 \max_j |u_j|^2 \le \alpha^2,$$

we determine that the Lipschitz constant $L \le \alpha^2$. According to Theorem B.1, the EMD algorithm ostensibly requires $\widetilde{\mathrm{O}}(\alpha^4)$ iterations to deliver a solution to (B.1) with constant precision. In practice, far fewer iterations suffice.

Remark B.2. *The application of EMD to* (B.1) *closely resembles the multiplicative weights method* [Kal07, Ch. 6] *for solving the* MAXCUT *problem* (3.4)*. Indeed, the two algorithms are substantially identical, except for the specific choice of step sizes and the method for constructing the final solution from the sequence of iterates. The efficiency estimates are also similar, except that the multiplicative weights method uses the widths of the constraints in lieu of the Lipschitz constant. EMD appears to be more effective in practice because it exploits the geometry of the problem more completely.*

B.5. **Grothendieck factorization via EMD.** Suppose $\boldsymbol{G}$ is an $s \times s$ Hermitian matrix. The Grothendieck factorization problem (5.2) can be expressed as solving

$$\min \ \lambda_{\max} \begin{bmatrix} -\alpha \operatorname{diag}(\boldsymbol{f}) & \boldsymbol{G} \\ \boldsymbol{G} & -\alpha \operatorname{diag}(\boldsymbol{f}) \end{bmatrix} \qquad \text{subject to} \qquad \boldsymbol{f} \in \Delta_s.$$

Abbreviate the objective function $J : \Delta_s \to \mathbb{R}$. Once again, EMD is an appropriate technique.

We may obtain subgradients using the same methods as before. Compute a normalized, maximal eigenvector of the matrix:

$$\begin{bmatrix} -\alpha \operatorname{diag}(\boldsymbol{f}) & \boldsymbol{G} \\ \boldsymbol{G} & -\alpha \operatorname{diag}(\boldsymbol{f}) \end{bmatrix} \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} = J(\boldsymbol{f}) \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \end{bmatrix} \quad \text{where} \quad \|\boldsymbol{u}\|_2^2 + \|\boldsymbol{v}\|_2^2 = 1.$$

Then a subgradient $\boldsymbol{\theta} \in \partial J(\boldsymbol{f})$ can be obtained from the formula

$$\boldsymbol{\theta} = -\alpha \left(|\boldsymbol{u}|^2 + |\boldsymbol{v}|^2 \right).$$

The Lipschitz constant $L \leq \alpha$, so the number of iterations of EMD is apparently $\widetilde{\mathrm{O}}(\alpha^2)$. Of course, the eigenvector calculations can be streamlined by exploiting the structure of the matrix.

Acknowledgments

The author thanks Ben Recht for helpful discussions about eigenvalue minimization.

References

[Ali95] F. Alizadeh. Interior-point methods in semidefinite programming with applications to combinatorial optimization. *SIAM J. Optimization*, 5(1):13–51, Feb. 1995.

[AN04] N. Alon and A. Naor. Approximating the cut norm via Grothendieck's inequality. In *Proc. 36th Ann. ACM Symposium on Theory of Computing (STOC)*, pages 72–80, Chicago, 2004.

[BDM08] C. Boutsidis, P. Drineas, and M. Mahoney. On selecting exactly k columns from a matrix. Submitted for publication, 2008.

[BT87] J. Bourgain and L. Tzafriri. Invertibility of "large" submatrices with applications to the geometry of Banach spaces and harmonic analysis. *Israel J. Math*, 57(2):137–224, 1987.

[BT91] J. Bourgain and L. Tzafriri. On a problem of Kadison and Singer. *J. reine angew. Math.*, 420:1–43, 1991.

[BT03] A. Beck and M. Teboulle. Mirror descent and nonlinear projected subgradient methods for convex optimization. *Operations Res. Lett.*, 31:167–175, 2003.

[GE96] M. Gu and S. Eisenstat. Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM J. Sci. Comput.*, 17(4):848–869, Jul. 1996.

[GVL96] G. H. Golub and C. F. Van Loan. *Matrix Computations*. Johns Hopkins Univ. Press, 3rd edition, 1996.

[GW95] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995.

[HJ85] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge Univ. Press, 1985.

[Kal07] S. Kale. *Efficient algorithms using the multiplicative weights update method*. Ph.D. dissertation, Computer Science Dept., Princeton Univ., Princeton, Nov. 2007.

[Lew96] A. S. Lewis. Convex analysis on the Hermitian matrices. *SIAM J. Optimization*, 6:164–177, 1996.

[LO96] A. S. Lewis and M. L. Overton. Eigenvalue optimization. *Acta Numerica*, 5:149–190, 1996.

[LT91] M. Ledoux and M. Talagrand. *Probability in Banach Spaces: Isoperimetry and Processes*. Springer, 1991.

[Pis86] G. Pisier. *Factorization of linear operators and geometry of Banach spaces*. Number 60 in CBMS Regional Conference Series in Mathematics. AMS, Providence, 1986. Reprinted with corrections, 1987.

[Roc70] R. T. Rockafellar. *Convex Analysis*. Princeton Univ. Press, 1970.

[Roh00] J. Rohn. Computing the norm $\|A\|_{\infty,1}$ is NP-hard. *Linear and Multilinear Algebra*, 47:195–204, 2000.

[RV07] M. Rudelson and R. Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *J. Amer. Comput. Soc.*, 54(4):Article 21, pp. 1–19, Jul. 2007.

[Sza90] S. Szarek. Spaces with large distance from ℓ_∞^n and random matrices. *Amer. J. Math.*, 112(6):899–942, Dec. 1990.

[Tro08] J. A. Tropp. On the linear independence of spikes and sines. *J. Fourier Anal. Appl.*, 2008. To appear.

[Ver01] R. Vershynin. Johns decompositions: Selecting a large part. *Israel J. Math.*, 122:253–277, 2001.

[Ver06] R. Vershynin. *High Dimensional Probability*, volume 51 of *IMS Lecture Notes—Monograph Series*, chapter Random sets of isomorphism of linear operators on Hilbert space, pages 148–154. IMS, 2006.