BINARY COMPONENT DECOMPOSITION
PART II: THE ASYMMETRIC CASE

RICHARD KUENG AND JOEL A. TROPP

ABSTRACT. This paper studies the problem of decomposing a low-rank matrix into a factor with binary entries, either from $\{\pm 1\}$ or from $\{0, 1\}$, and an unconstrained factor. The research answers fundamental questions about the existence and uniqueness of these decompositions. It also leads to tractable factorization algorithms that succeed under a mild deterministic condition. This work builds on a companion paper that addresses the related problem of decomposing a low-rank positive-semidefinite matrix into symmetric binary factors.

1. MOTIVATION

Constrained matrix decompositions are among the basic methods for unsupervised data analysis. These techniques play a role in many scientific and engineering fields, ranging from environmental engineering [PT94] and neuroscience [OF96] to signal processing [Com94] and statistics [ZHT06]. Constrained factorizations are powerful tools for identifying latent structure in a matrix; they also support data compression, summarization, and visualization.

The literature contains a number of frameworks [TB99, CDS02, Tro04, Sre04, Wit10, Jag11, Bac13, Ude15, BE16, Bru17, HV19] for thinking about constrained matrix factorization and for developing algorithms that pursue these factorizations. Nevertheless, we still lack theory that fully justifies these approaches. For instance, researchers have only attained a partial understanding of which factorization models are identifiable and which ones we can compute provably using efficient algorithms.

The purpose of this paper and its companion [KT19] is to develop foundational results on factorization models that we call *binary component decompositions*. In these models, one (or both) of the factors takes values in the set $\{\pm 1\}$ or in the set $\{0, 1\}$. Binary component decompositions are appropriate when the latent factors reflect an exclusive choice. From a mathematical perspective, these constrained factorizations also happen to be among the easiest ones to understand.

In this second paper, we consider the problem of factorizing a rectangular matrix into a binary factor and an unconstrained matrix of weights. We develop results on existence, uniqueness, tractable computation, and robustness to gross errors. Our analysis builds heavily on the work in the companion paper [KT19], which treats the problem of decomposing a positive-semidefinite matrix into symmetric binary factors.

1.1. **Notation.** We rely on standard notation from linear algebra and optimization. Scalars are written with lowercase Roman or Greek letters (x, ξ); lowercase bold letters $(\boldsymbol{x}, \boldsymbol{\xi})$ denote (column) vectors; uppercase bold letters $(\boldsymbol{X}, \boldsymbol{\Xi})$ denote matrices. We reserve calligraphic letters $(\mathcal{X})$ for sets. The symbol $\lesssim$ suppresses universal constants.

Throughout, n and m are natural numbers. We work in the real linear spaces $\mathbb{R}^n$ and $\mathbb{R}^m$ equipped with the standard inner product $\langle \cdot, \cdot \rangle$ and the associated norm topology featuring $\|\boldsymbol{x}\|_{\ell_2} = \sqrt{\langle \boldsymbol{x}, \boldsymbol{x} \rangle}$. The standard basis vector $\mathbf{e}_i$ has a one in the ith coordinate and zeros elsewhere, while $\mathbf{e}$ is the vector of ones; the dimension of these vectors depends on context. The map $^{\mathrm{t}}$ transposes a vector or matrix. The

binary operator $\odot$ is the Schur (i.e., componentwise) product of vectors. The closed and open probability simplices are the sets

$$\Delta_r = \left\{ \boldsymbol{\tau} \in \mathbb{R}^r : \tau_i \geq 0 \text{ and } \sum_{i=1}^r \tau_i = 1 \right\} \quad \text{and} \quad \Delta_r^+ = \left\{ \boldsymbol{\tau} \in \mathbb{R}^r : \tau_i > 0 \text{ and } \sum_{i=1}^r \tau_i = 1 \right\}.$$

We write $\mathbb{H}_n$ for the linear space of symmetric $n \times n$ real matrices. The symbol $\mathbf{I}$ denotes the identity matrix, and $\mathbf{E}$ denotes the matrix of ones; their dimensions are determined by the context. The dagger † refers to the Moore–Penrose pseudoinverse. A positive-semidefinite (psd) matrix is a symmetric matrix $\boldsymbol{X}$ that satisfies $\boldsymbol{u}^{\mathrm{t}} \boldsymbol{X} \boldsymbol{u} \geq 0$ for all vectors $\boldsymbol{u}$ with compatible dimension. The statement $\boldsymbol{X} \succcurlyeq \boldsymbol{Y}$ means that $\boldsymbol{X} - \boldsymbol{Y}$ is psd, and $\boldsymbol{X} \succ \boldsymbol{Y}$ means that $\boldsymbol{X} - \boldsymbol{Y}$ is strictly positive definite, i.e. $\boldsymbol{u}^{\mathrm{t}} (\boldsymbol{X} - \boldsymbol{Y}) \boldsymbol{u} > 0$ for all vectors $\boldsymbol{u}$ with compatible dimension.

2. SIGN COMPONENT DECOMPOSITION AND BINARY COMPONENT DECOMPOSITION

We begin with a short discussion of the singular-value decomposition and its properties (Section 2.1). Afterward, we introduce the two factorizations that we treat in this paper, the sign component decomposition (Section 2.2) and the binary component decomposition (Section 2.3). We present our main results on situations where these factorizations are uniquely determined and when they can be computed using efficient algorithms. An outline of the rest of the paper appears in Section 2.5.

2.1. **The singular-value decomposition.** We begin with the singular-value decomposition (SVD), the royal emperor among all matrix factorizations. Let $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ be a rectangular matrix. For some natural number $r \leq \min\{m, n\}$, we can decompose this matrix as

$$\boldsymbol{B} = \sum_{i=1}^r \sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^{\mathrm{t}}. \tag{2.1}$$

In this expression, $\{\boldsymbol{u}_1, \dots, \boldsymbol{u}_r\} \subset \mathbb{R}^n$ and $\{\boldsymbol{v}_1, \dots, \boldsymbol{v}_r\} \subset \mathbb{R}^m$ are orthonormal families of left and right singular vectors associated with the positive singular values $\sigma_1 \geq \sigma_2 \geq \cdots \geq \sigma_r > 0$. We can also convert the decomposition (2.1) into a matrix factorization:

$$\boldsymbol{B} = \boldsymbol{U} \boldsymbol{\Sigma} \boldsymbol{V}^{\mathrm{t}} \quad \text{where} \quad \begin{aligned} \boldsymbol{U} &= \begin{bmatrix} \boldsymbol{u}_1 & \dots & \boldsymbol{u}_r \end{bmatrix} \in \mathbb{R}^{n \times r}; \\ \boldsymbol{\Sigma} &= \mathrm{diag}(\sigma_1, \dots, \sigma_r) \in \mathbb{R}^{r \times r}; \\ \boldsymbol{V} &= \begin{bmatrix} \boldsymbol{v}_1 & \dots & \boldsymbol{v}_r \end{bmatrix} \in \mathbb{R}^{m \times r}. \end{aligned} \tag{2.2}$$

The matrices $\boldsymbol{U}$ and $\boldsymbol{V}$ are orthonormal; that is, $\boldsymbol{U}^{\mathrm{t}} \boldsymbol{U} = \mathbf{I}$ and $\boldsymbol{V}^{\mathrm{t}} \boldsymbol{V} = \mathbf{I}$.

The singular-value decomposition is intimately connected to the problem of finding a best low-rank approximation of a matrix [Mir60]. Indeed, for any unitarily invariant norm $\|\cdot\|$,

$$\min_{\mathrm{rank}\, \boldsymbol{L} = k} \|\boldsymbol{B} - \boldsymbol{L}\| = \left\| \boldsymbol{B} - \sum_{i=1}^k \sigma_i \boldsymbol{u}_i \boldsymbol{v}_i^{\mathrm{t}} \right\| \quad \text{for each } k = 1, \dots, r.$$

This variational property has a wide range of consequences, both theoretical and applied.

The singular-value decomposition also holds a distinguished place in statistics because of its connection with principal component analysis [Jol02]. Given a data matrix $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ with standardized[1] rows, we can perform a singular-value decomposition to express $\boldsymbol{B} = \boldsymbol{U} \boldsymbol{W}^{\mathrm{t}}$, where $\boldsymbol{W} = \boldsymbol{\Sigma} \boldsymbol{V}^{\mathrm{t}}$. In this setting, the left singular vectors $\boldsymbol{u}_i$ are called *principal components*, the directions in which the columns of $\boldsymbol{B}$ exhibit the most variability. The entries of the matrix $\boldsymbol{W}$ are called *weights* or *loadings*; they are the coefficients with which we combine the principal components to express the original data points.

On the positive side of the ledger, the singular-value decomposition (2.1)–(2.2) always exists, and it is uniquely determined when the (nonzero) singular values are distinct. Moreover, we can compute the singular-value decomposition, up to a fixed (high) accuracy, by means of highly refined algorithms, in polynomial time.

On the negative side, we cannot impose constraints on the singular vectors to enforce prior knowledge about the data. Second, we generally cannot assign an interpretation or meaning to the singular vectors, without committing the sin of reification. Moreover, the orthogonality of singular vectors may not be an

---

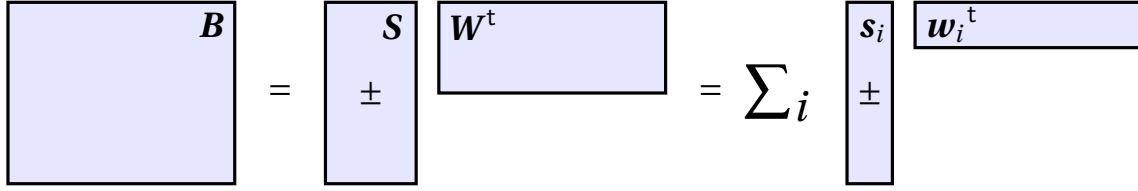[1]A vector is *standardized* if its entries sum to zero and its Euclidean norm equals one.

FIGURE 2.1: *Asymmetric sign component decomposition.* The sign component decomposition (2.3)–(2.4) expresses a rectangular matrix B as the product of a sign matrix S and an unconstrained matrix W^{t}.

appropriate constraint in applications. Structured matrix factorizations are designed to address one or more of these shortcomings.

2.2. **Sign component decomposition.** In this project, we consider matrix factorization models where one of the factors is required to take binary values. In this section, we treat the case where the entries of the binary factor are limited to the set $\{\pm 1\}$. In Section 2.3, we turn to the case where the entries are drawn from the set $\{0, 1\}$.

2.2.1. *The decomposition.* As before, assume that $B \in \mathbb{R}^{n \times m}$ is a rectangular matrix. We seek a decomposition of the form

$$B = SW^{\mathsf{t}} \quad \text{where} \quad S \in \{\pm 1\}^{n \times r} \quad \text{and} \quad W \in \mathbb{R}^{m \times r}. \tag{2.3}$$

This factorization can also be written in vector notation as

$$B = \sum_{i=1}^{r} s_i w_i^{\mathsf{t}} \quad \text{where} \quad s_i \in \{\pm 1\}^n \quad \text{and} \quad w_i \in \mathbb{R}^m. \tag{2.4}$$

We call (2.3)–(2.4) an (asymmetric) *sign component decomposition* of the matrix B. The left factor S is called the *sign component*; its columns s_i are also called *sign components*. The right factor W is unconstrained; its entries are called *weights* or *loadings*. See Figure 2.1 for an illustration.

It is not hard to show that each $n \times m$ matrix B admits a plethora of distinct sign component decompositions (2.3) where the inner dimension is $r = n$; see Proposition 4.1. It is more interesting to consider a low-rank matrix B and to search for *minimal* decompositions, those where the inner dimension r of the factorization (2.3) equals the rank of B.

Remark 2.1 (Matrix sign function)**.** The sign component decomposition must not be confused with the matrix sign function, which is a spectral computation related to the polar factorization [Hig08, Chap. 5].

2.2.2. *Schur independence.* The sign component decomposition (2.3)–(2.4) has a combinatorial quality, which suggests that it might be hard to find. Remarkably, there is a large class of matrices for which we can tractably compute a minimal sign component decomposition. The core requirement is that the sign components must be somewhat different. The following definition [LP96, KT19] encapsulates this idea.

Definition 2.2 (Schur independence of sign vectors)**.** A set $\{s_1, \ldots, s_r\} \subseteq \{\pm 1\}^n$ of sign vectors is *Schur independent* when the set

$$\{e\} \cup \{s_i \odot s_j : 1 \le i < j \le r\} \subseteq \mathbb{R}^n \quad \text{is linearly independent.}$$

By extension, we also say that the sign matrix $S = \begin{bmatrix} s_1 & \ldots & s_r \end{bmatrix} \in \{\pm 1\}^{n \times r}$ is Schur independent when its columns form a Schur independent set.

Let us summarize the basic properties of Schur independent sets [LP96, Tro18, KT19].

Fact 2.3 (Schur independence)**.** *Assume that the set $\mathscr{S} = \{s_1, \ldots, s_r\} \subseteq \{\pm 1\}^n$ of sign vectors is Schur independent. We have the following consequences.*

 (1) The family $\mathscr{S}$ is linearly independent.
 (2) Each subset of $\mathscr{S}$ is Schur independent.
 (3) For any choice $\xi \in \{\pm 1\}^r$ of signs, the set $\{\xi_1 s_1, \ldots, \xi_r s_r\}$ remains Schur independent.

(4) The cardinality r of the set $\mathscr{S}$ satisfies $r \leq \frac{1}{2}(1 + \sqrt{8n - 7})$.
(5) We can determine whether or not $\mathscr{S}$ is Schur independent in polynomial time.

Schur independence is best understood as a kind of "general position" property for sign vectors. Roughly speaking, almost all collections of sign vectors are Schur independent, provided that the cardinality r meets the bound stated in Fact 2.3(4). This intuition is quantified in the paper [Tro18].

2.2.3. *Computation.* The main result of this paper is an algorithm for computing the minimal asymmetric sign component decomposition of a low-rank matrix. This algorithm succeeds precisely when the sign component is Schur independent. Moreover, this condition is sufficient to ensure that the sign component decomposition is essentially unique.

Theorem I (Sign component decomposition)**.** *Let $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ be a matrix that admits a sign component decomposition $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{W}^{\mathrm{t}}$ where*

(1) The sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ is Schur independent;
(2) The weight matrix $\boldsymbol{W} \in \mathbb{R}^{m \times r}$ has full column rank.

Then the minimal sign component decomposition (with inner dimension r) is determined up to simultaneous sign flips and permutations of the columns of the factors. Algorithm 1 computes this decomposition in time polynomial in $n + m$.

The uniqueness claim is a consequence of Theorem 4.4, while the computational claim follows from Theorem 5.1.

Theorem I identifies a rich set of factorizable matrices for which exact identification is always tractable and essentially unique. Moreover, existing denoising techniques allow us to compute the factorization in the presence of gross errors; see Section 7. Small perturbations appear more challenging; we will study this problem in future work.

It is surprising that the exact sign component decomposition is tractable. Most existing approaches to structured matrix factorization only produce approximations, and many of these approaches lack rigorous guarantees. The companion paper [KT19, Sec. 8] contains a discussion of the related work.

2.3. **Binary component decomposition.** The asymmetric sign component decomposition also serves as a primitive that allows us to compute other discrete matrix factorizations. In this section, we turn to the problem of producing a decomposition where one component takes values in the set $\{0, 1\}$.

2.3.1. *The decomposition.* Suppose that $\boldsymbol{C} \in \mathbb{R}^{n \times m}$ is a rectangular matrix. We consider a decomposition of the form

$$\boldsymbol{C} = \boldsymbol{Z}\boldsymbol{W}^{\mathrm{t}} \quad \text{where} \quad \boldsymbol{Z} \in \{0, 1\}^{n \times r} \quad \text{and} \quad \boldsymbol{W} \in \mathbb{R}^{m \times r}. \tag{2.5}$$

The vector formulation of this decomposition is

$$\boldsymbol{C} = \sum_{i=1}^{r} \boldsymbol{z}_i \boldsymbol{w}_i^{\mathrm{t}} \quad \text{where} \quad \boldsymbol{z}_i \in \{0, 1\}^n \quad \text{and} \quad \boldsymbol{w}_i \in \mathbb{R}^m. \tag{2.6}$$

We refer to (2.5)–(2.6) as an (asymmetric) *binary component decomposition* of the matrix $\boldsymbol{C}$. The left factor $\boldsymbol{Z}$ is called the *binary component*, and its columns $\boldsymbol{z}_i$ are also called *binary components*. The right factor $\boldsymbol{W}$ is unconstrained; we refer to it as a *weight* matrix.

Every $n \times m$ matrix $\boldsymbol{C}$ admits a superabundance of distinct binary component decompositions (2.5) where the inner dimension $r = n$. We focus on the case where the matrix $\boldsymbol{C}$ has low rank, and the factorization is *minimal*; that is, the inner dimension r in (2.5) equals the rank of $\boldsymbol{C}$.

2.3.2. *Schur independence.* We can reduce the problem of computing a binary component decomposition to the problem of computing a sign component decomposition.

To do so, we first observe that there is an affine map that places the binary vectors and sign vectors in one-to-one correspondence:

$$\boldsymbol{F} : \{0, 1\}^n \to \{\pm 1\}^n \quad \text{where} \quad \boldsymbol{F} : \boldsymbol{z} \mapsto 2\boldsymbol{z} - \mathbf{e} \quad \text{and} \quad \boldsymbol{F}^{-1} : \boldsymbol{s} \mapsto \frac{1}{2}(\boldsymbol{s} + \mathbf{e}). \tag{2.7}$$

Algorithm 1 *Asymmetric sign component decomposition* (2.3)–(2.4) *of a matrix with a Schur independent sign component.* Implements the procedure from Section 5.

Input: Rank-r matrix $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ that satisfies the conditions of Theorem I.
Output: Sign matrix $\tilde{\boldsymbol{S}} \in \{\pm 1\}^{n \times r}$ and weight matrix $\tilde{\boldsymbol{W}} \in \mathbb{R}^{m \times r}$ where $\boldsymbol{B} = \tilde{\boldsymbol{S}} \tilde{\boldsymbol{W}}^{\mathrm{t}}$.

1 **function** ASYMSIGNCOMPONENTDECOMPOSITION($\boldsymbol{B}$)
2     $[n, m] \leftarrow \mathtt{size}(\boldsymbol{B})$ and $r \leftarrow \mathrm{rank}(\boldsymbol{A})$
3     $\boldsymbol{U} \leftarrow \mathtt{orth}(\boldsymbol{A})$ ▷ Find a basis for the range of $\boldsymbol{A}$
4   Find the solution $(\boldsymbol{X}_\star, \boldsymbol{Y}_\star)$ to the semidefinite program

$$\underset{\boldsymbol{X} \in \mathbb{H}_n, \boldsymbol{Y} \in \mathbb{H}_m}{\text{minimize}} \quad \mathrm{trace}(\boldsymbol{Y})$$

$$\text{subject to} \quad \mathrm{trace}(\boldsymbol{U}^{\mathrm{t}} \boldsymbol{X} \boldsymbol{U}) = n \text{ and } \mathrm{diag}(\boldsymbol{X}) = \mathbf{e};$$

$$\begin{bmatrix} \boldsymbol{X} & \boldsymbol{B} \\ \boldsymbol{B}^{\mathrm{t}} & \boldsymbol{Y} \end{bmatrix} \succcurlyeq \mathbf{0}$$

5   Apply Algorithm 3 to $\boldsymbol{X}_\star$ to obtain a symmetric sign component decomposition (3.1):

$$\boldsymbol{X}_\star = \tilde{\boldsymbol{S}} \, \mathrm{diag}(\tilde{\boldsymbol{\tau}}) \, \tilde{\boldsymbol{S}}^{\mathrm{t}} \quad \text{where } \tilde{\boldsymbol{S}} \in \{\pm 1\}^{n \times r}$$

6   Find the solution $\tilde{\boldsymbol{W}} \in \mathbb{R}^{m \times r}$ to the linear system $\boldsymbol{B} = \tilde{\boldsymbol{S}} \tilde{\boldsymbol{W}}^{\mathrm{t}}$

We can extend the map $\boldsymbol{F}$ to a matrix by applying it to each column. This correspondence suggests that there should also be a concept of Schur independence for binary vectors. Here is the notion that suits our purposes.

Definition 2.4 (Schur independence of binary vectors). A set $\{\boldsymbol{z}_1, \dots, \boldsymbol{z}_r\} \subseteq \{0, 1\}^n$ of binary vectors is *Schur independent* when the set

$$\{\mathbf{e}\} \cup \{\boldsymbol{z}_i : 1 \le i \le r\} \cup \{\boldsymbol{z}_i \odot \boldsymbol{z}_j : 1 \le i < j \le r\} \subseteq \mathbb{R}^n \quad \text{is linearly independent.}$$

By extension, we say that a binary matrix $\boldsymbol{Z} = \begin{bmatrix} \boldsymbol{z}_1 & \dots & \boldsymbol{z}_r \end{bmatrix} \in \{0, 1\}^{n \times r}$ is Schur independent when its columns compose a Schur independent set.

The following result [KT19, Prop. 6.3] describes the precise connection between the two flavors of Schur independence.

Fact 2.5 (Kueng & Tropp). *The binary matrix $\boldsymbol{Z} \in \{0, 1\}^r$ is Schur independent if and only if the sign matrix $\begin{bmatrix} \boldsymbol{F}(\boldsymbol{Z}) & \mathbf{e} \end{bmatrix} \in \{\pm 1\}^{r+1}$ is Schur independent.*

2.3.3. *Computation.* With these definitions at hand, we can state our main result on binary component decompositions.

Theorem II (Binary component decomposition). *Let $\boldsymbol{C} \in \mathbb{R}^{n \times m}$ be a matrix that admits a binary component decomposition $\boldsymbol{C} = \boldsymbol{Z} \boldsymbol{W}^{\mathrm{t}}$ where*

    (1) The binary matrix $\boldsymbol{Z} \in \{0, 1\}^{n \times r}$ is Schur independent;
    (2) The weight matrix $\boldsymbol{W} \in \mathbb{R}^{m \times r}$ has full column rank.

Then the minimal binary component decomposition (with inner dimension r) is determined up to simultaneous permutation of the columns of the factors. Algorithm 2 computes the decomposition in time polynomial in $n + m$.

The uniqueness claim is established in Theorem 6.4, and the computational claim appears in Theorem 6.5.

2.4. **The planted sign basis problem.** Theorem I and Theorem II allow us to solve some interesting combinatorial problems in linear algebra.

Algorithm 2 *Binary component decomposition* (2.5) *of a matrix with a Schur independent binary component.* Implements the procedure from Section 6.5.

Input: Rank-r matrix $C \in \mathbb{H}_n$ that satisfies the conditions of Theorem II
Output: Binary matrix $\tilde{Z} \in \{0,1\}^{n \times r}$ and weight matrix $\tilde{W}_+ \in \mathbb{R}^{m \times r}$ for which $C = \tilde{Z}\tilde{W}_+^{\mathsf{t}}$.

1 **function** ASYMBINARYCOMPONENTDECOMPOSITION(C)
2   Form the matrix $B = 2C - \mathbf{E}$
3   Apply Algorithm 1 to B to obtain a sign component decomposition
$$B = \tilde{S}\tilde{W}^{\mathsf{t}} = \begin{bmatrix} \tilde{s}_1 & \dots & \tilde{s}_{r+1} \end{bmatrix} \begin{bmatrix} \tilde{w}_1 & \dots & \tilde{w}_{r+1} \end{bmatrix}^{\mathsf{t}}$$
4   Find the index i and sign $\varphi \in \{\pm 1\}$ where $\tilde{s}_i = \varphi \mathbf{e}$. Permute the sequences to interchange
$$\tilde{s}_i \leftrightarrow \tilde{s}_{r+1} \quad \text{and} \quad \tilde{w}_i \leftrightarrow \tilde{w}_{r+1}$$
5   Find the solution $\xi \in \mathbb{R}^r$ to the linear system
$$\begin{bmatrix} \tilde{w}_1 & \dots & \tilde{w}_r \end{bmatrix} \xi = \varphi \tilde{w}_{r+1} + \mathbf{e}$$
6   Set $\tilde{z}_i = \frac{1}{2}(\xi_i \tilde{s}_i + \mathbf{e})$ for each index $i = 1, \dots, r$
7   Define the binary matrix $\tilde{Z} = \begin{bmatrix} \tilde{z}_1 & \dots & \tilde{z}_r \end{bmatrix}$ and the weight matrix $\tilde{W}_+ = \begin{bmatrix} \xi_1 \tilde{w}_1 & \dots & \xi_r \tilde{w}_r \end{bmatrix}$

Problem 2.6 (Planted sign basis). Let $\mathsf{L} \subseteq \mathbb{R}^n$ be an r-dimensional subspace that admits a sign basis:

$$\mathsf{L} = \mathrm{span}\{s_1, \dots, s_r\} \quad \text{where each } s_i \in \{\pm 1\}^n.$$

Given the subspace L, find a sign basis for the subspace.

To clarify, we can assume that the problem data is a matrix $B \in \mathbb{R}^{n \times m}$ whose range equals the r-dimensional subspace L. We must output a set of r sign vectors that generates the subspace. The brute force approach may require us to sift through around 2^{nr} families of sign vectors. Is it possible to solve the problem more efficiently?

Let us outline a solution for Problem 2.6 in the case where L has a sign basis $\{s_1, \dots, s_r\} \subseteq \{\pm 1\}^n$ that is Schur independent. This is a rather mild deterministic condition, provided that the dimension r of the subspace satisfies $r < \frac{1}{2}(1 + \sqrt{8n-7})$. The hypothesis also guarantees that the basis is determined up to permutation and sign flips, per Theorem I.

Here is how we solve the problem. Let $B \in \mathbb{R}^{n \times m}$ be a matrix whose range coincides with the subspace L. A Schur independent set is linearly independent, so we can write the matrix in the form $B = SW^{\mathsf{t}}$, where $S = \begin{bmatrix} s_1 & \dots & s_r \end{bmatrix} \in \{\pm 1\}^{n \times r}$ and the weight matrix $W \in \mathbb{R}^{m \times r}$ has full column rank. As a consequence, we can apply Algorithm 1 to the matrix B to obtain a sign component decomposition $B = \tilde{S}\tilde{W}^{\mathsf{t}}$. Theorem I ensures that the columns of $\tilde{S}$ coincide with the columns of S up to sign flips and permutations. In other words, the columns of $\tilde{S}$ compose the (unique) sign basis that generates L. In summary, we can solve Problem 2.6 for any subspace that is spanned by a Schur independent family of sign vectors.

A similar procedure, using Algorithm 2, allows us to solve a variant of Problem 2.6 where we seek a planted *binary* basis for a subspace. Indeed, if a subspace is generated by a Schur independent family of binary vectors, then we can identify the basis up to permutation.

2.5. **Roadmap.** We continue with a discussion about symmetric sign component decompositions in Section 3. In Section 4, we develop basic results about existence and uniqueness of asymmetric sign component decompositions. Section 5 explains how to compute an sign component decomposition. We turn to binary component decomposition in Section 6. Finally, in Section 7, we state some results on robustness of sign component decomposition which we prove in the appendices. For a discussion of related work, see the companion paper [KT19, Sec. 8].

3. SYMMETRIC SIGN COMPONENT DECOMPOSITION

This section contains a summary of the principal results from the companion paper [KT19]. These results play a core role in our study of asymmetric factorizations.

3.1. **Signed permutations.** Matrix factorizations are usually not fully determined because they are invariant under some group of symmetries. For example, consider the decomposition of a psd matrix as the outer product of two symmetric factors:

$$A = BB^{\mathrm{t}} = (BQ)(BQ)^{\mathrm{t}} \quad \text{for each orthogonal } Q.$$

Each of the factorizations on the right is equally valid, because there is no constraint that forbids rotations.

For binary component decompositions, permutations compose the relevant symmetry group.

Definition 3.1 (Permutation). A *permutation* on r letters is an element π of the symmetric group Sym_r. A permutation π acts on $\mathbb{R}^r$ via the linear map $x \mapsto (x_{\pi(1)}, \ldots, x_{\pi(r)})$. This linear map can be represented by the *permutation matrix* $\Pi \in \mathbb{R}^{r \times r}$ whose entries take the form $(\Pi)_{ij} = 1$ where $j = \pi(i)$ and are zero otherwise. A permutation matrix is orthogonal: $\Pi\Pi^{\mathrm{t}} = I = \Pi^{\mathrm{t}}\Pi$.

For sign component decompositions, the signed permutations make up the relevant symmetry group.

Definition 3.2 (Signed permutation). A *signed permutation* on r letters is a pair $(\pi, \boldsymbol{\xi}) \in \mathsf{Sym}_r \times \{\pm 1\}^r$ consisting of a permutation π on r letters and a sign vector $\boldsymbol{\xi} \in \{\pm 1\}^r$. The signed permutation $(\pi, \boldsymbol{\xi})$ acts on $\mathbb{R}^r$ via the linear map $x \mapsto (\xi_1 x_{\pi(1)}, \ldots, \xi_r x_{\pi(r)})$. This linear map can also be represented by the *signed permutation matrix* $\Pi \in \mathbb{R}^{r \times r}$ whose entries satisfy $(\Pi)_{ij} = \xi_i$ when $j = \pi(i)$ and are otherwise zero. Each signed permutation matrix is orthogonal.

3.2. **Symmetric sign component decomposition.** In the companion paper [KT19], we explored the problem of computing a (symmetric) sign component decomposition of a correlation matrix. This research provides the foundation for the asymmetric sign component decomposition. Let us take a moment to present the principal definitions and results from the associated work.

Let $A \in \mathbb{H}_n$ be a correlation matrix; that is, A is psd with all diagonal entries equal to one. We say that A has a *symmetric sign component decomposition* when

$$A = S \operatorname{diag}(\boldsymbol{\tau}) S^{\mathrm{t}} \quad \text{where} \quad S \in \{\pm 1\}^{n \times r} \text{ and } \boldsymbol{\tau} \in \Delta_r^+. \tag{3.1}$$

In vector form,

$$A = \sum_{i=1}^{r} \tau_i s_i s_i^{\mathrm{t}} \quad \text{where} \quad s_i \in \{\pm 1\}^n \quad \text{and} \quad (\tau_1, \ldots, \tau_r) \in \Delta_r^+.$$

The sign matrix S is called the *sign component*, while the positive diagonal matrix, $\operatorname{diag}(\boldsymbol{\tau})$, is a list of convex coefficients. Not all correlation matrices admit a symmetric sign component decomposition, nor does the factorization need to be uniquely determined; see [KT19] for a full discussion.

The situation improves markedly when the sign component S is Schur independent. In this case, the sign component decomposition is essentially unique, and we can compute it by means of an efficient algorithm [KT19, Thm. I].

Fact 3.3 (Kueng & Tropp). *Let $A \in \mathbb{H}_n$ be a correlation matrix that admits a sign component decomposition:*

$$A = S \operatorname{diag}(\boldsymbol{\tau}) S^{\mathrm{t}} \quad \text{where} \quad S \in \{\pm 1\}^{n \times r} \text{ is Schur independent and } \boldsymbol{\tau} \in \Delta_r^+.$$

Then the sign component decomposition of A is determined up to signed permutation. Moreover, with probability one, Algorithm 3 computes the sign component decomposition. That is, the output is a pair $(\tilde{S}, \tilde{\boldsymbol{\tau}})$ where the sign matrix $\tilde{S} = S\Pi$ and the convex coefficients $\tilde{\tau}_i = |(\Pi^{\mathrm{t}}\boldsymbol{\tau})_i|$ $(1 \le i \le r)$, for a signed permutation matrix $\Pi \in \mathbb{R}^{r \times r}$.

A major ingredient in the proof of Fact 3.3 is a characterization of the set of correlation matrices that are generated by a Schur independent family of sign vectors [KT19, Thm. 3.6].

Algorithm 3 *Symmetric sign component decomposition* (3.1) *of a correlation matrix with a Schur independent sign component.* Duplicates [KT19, Alg. 1].

Input: Rank-r correlation matrix $A \in \mathbb{H}_n$ that satisfies the conditions of Fact 3.3.
Output: Sign matrix $\tilde{S} \in \{\pm 1\}^{n \times r}$ and convex coefficients $\tilde{\tau} \in \Delta_r^+$ where $A = \tilde{S} \operatorname{diag}(\tilde{\tau}) \tilde{S}^{\mathrm{t}}$.

1 **function** SIGNCOMPONENTDECOMPOSITION(A)
2      $[n, \sim] \leftarrow \texttt{size}(A)$ and $r \leftarrow \operatorname{rank}(A)$
3   **for** $i = 1$ to $(r - 1)$ **do**
4          $U \leftarrow \texttt{orth}(A)$ ▷ Find a basis for the range of A
5          $g \leftarrow \texttt{randn}(n, 1)$ ▷ Draw a random direction
6   Find the solution $X_\star$ to the semidefinite program

$$\underset{X \in \mathbb{S}^n}{\operatorname{maximize}} \quad g^{\mathrm{t}} X g \quad \text{subject to} \quad \operatorname{trace}\left(U^{\mathrm{t}} X U\right) = n \text{ and } \operatorname{diag}(X) = e \text{ and } X \succcurlyeq 0$$

7   Factorize the rank-one matrix $X_\star = \tilde{s}_i \tilde{s}_i^{\mathrm{t}}$ ▷ Extract a sign component
8   Find the solution $\zeta_\star$ to the semidefinite program

$$\underset{\zeta \in \mathbb{R}}{\operatorname{maximize}} \quad \zeta \quad \text{subject to} \quad \zeta A + (1 - \zeta) X_\star \succcurlyeq 0$$

9          $A \leftarrow \zeta_\star A + (1 - \zeta_\star) X_\star$
10   Factorize the rank-one matrix $A = \tilde{s}_r \tilde{s}_r^{\mathrm{t}}$ ▷ $\operatorname{rank}(A) = 1$ in final iteration
11   Define the matrix $\tilde{S} = \begin{bmatrix} \tilde{s}_1 & \dots & \tilde{s}_r \end{bmatrix}$, and find the solution $\tilde{\tau} \in \Delta_r^+$ to the linear system

$$A = \tilde{S} \operatorname{diag}(\tilde{\tau}) \tilde{S}^{\mathrm{t}}$$

Fact 3.4 (Kueng & Tropp). *Suppose that $S \in \{\pm 1\}^{n \times r}$ is a Schur independent sign matrix, and let $P \in \mathbb{H}_n$ be the orthogonal projector onto* range(S). *Then*

$$\{S \operatorname{diag}(\tau) S^{\mathrm{t}} : \tau \in \Delta_r\} = \{X \in \mathbb{H}_n : \operatorname{trace}(PX) = n \text{ and } \operatorname{diag}(X) = e \text{ and } X \succcurlyeq 0\}. \tag{3.2}$$

Fact 3.4 is a powerful tool for working with sign component decompositions. Indeed, we can compute the projector P onto the range of a Schur independent sign matrix S directly from any particular correlation matrix $A = S \operatorname{diag}(\tau) S^{\mathrm{t}}$ with $\tau \in \Delta_r^+$. As a consequence, the identity (3.2) provides an alternative representation for the set of all correlation matrices with sign component S, which allows us to optimize over this set. Fact 3.4 also plays a critical role in our method for computing an asymmetric sign component decomposition.

4. EXISTENCE AND UNIQUENESS OF THE ASYMMETRIC SIGN COMPONENT DECOMPOSITION

In this section, we begin our investigation of the asymmetric sign component decomposition. We lay out some of the basic questions, and we start to deliver the answers.

4.1. **Questions.** This paper addresses four fundamental problems raised by the definition (2.3)–(2.4) of the asymmetric sign component decomposition:

(1) **Existence:** Which matrices admit a sign component decomposition?
(2) **Uniqueness:** When is the sign component decomposition unique, modulo symmetries?
(3) **Computation:** How can we find a sign component decomposition in polynomial time?
(4) **Robustness:** How can we find a sign component decomposition from a noisy observation?

This section treats the structural questions about existence and uniqueness of the sign component decomposition, and Section 5 explains how we can compute the factorization. Last, Section 7 describes some situations where we can extract a sign component decomposition from imperfect data.

4.2. **Existence.** We quickly dispatch the first question, which concerns the existence of asymmetric sign component decompositions.

Proposition 4.1 (Sign component decomposition: Existence)**.** *Every matrix $B \in \mathbb{R}^{n \times m}$ admits a sign component decomposition* (2.3) *with inner dimension $r = n$.*

Proof. Let $S \in \{\pm 1\}^{n \times n}$ be a nonsingular matrix of signs. Define the second factor $W^{\mathrm{t}} = S^{\dagger} B$. $\qquad \square$

As an aside, we remark that nonsingular sign matrices are ubiquitous. Indeed, a uniformly random element of $\{\pm 1\}^{n \times n}$ is nonsingular with exceedingly high probability [Tik18].

Proposition 4.1 ensures that every matrix has an exorbitant number of sign component decompositions. Therefore, we need to burden the factorization with extra conditions before it is determined uniquely. We intend to focus on minimal factorizations, where the target matrix $B \in \mathbb{R}^{n \times m}$ has rank $r < n$, and the number r of sign components coincides with the rank.

4.3. **Symmetries.** Like many other matrix factorizations, the sign component decomposition has some symmetries that we can never resolve. Before we can turn to the question of uniqueness, we need to discuss invariants of the factorization.

Signed permutations preserve the sign component decomposition (2.3)–(2.4) in the following sense. Suppose that $B \in \mathbb{R}^{n \times m}$ has the sign component decomposition

$$B = SW^{\mathrm{t}} = \begin{bmatrix} s_1 & \dots & s_r \end{bmatrix} \begin{bmatrix} w_1 & \dots & w_r \end{bmatrix}^{\mathrm{t}} = \sum_{i=1}^{r} s_i w_i^{\mathrm{t}}.$$

For a signed permutation $(\pi, \boldsymbol{\xi})$ on r letters with associated signed permutation matrix Π, we have

$$B = S\Pi\Pi^{-1}W^{\mathrm{t}} = (S\Pi)(W\Pi)^{\mathrm{t}} = \begin{bmatrix} \xi_1 s_{\pi(1)} & \dots & \xi_r s_{\pi(r)} \end{bmatrix} \begin{bmatrix} \xi_1 w_{\pi(1)} & \dots & \xi_r w_{\pi(r)} \end{bmatrix}^{\mathrm{t}}.$$

Observe that $S\Pi \in \{\pm 1\}^{n \times r}$ remains a sign matrix. Therefore, SW^{t} and $(S\Pi)(W\Pi)^{\mathrm{t}}$ are both sign component decompositions of B.

We have no cause to prefer one of the sign component decompositions induced by a signed permutation over the others. Thus, it is appropriate to treat them all as equivalent.

Definition 4.2 (Sign component decomposition: Equivalence)**.** Suppose that $B = SW^{\mathrm{t}}$ and $B = \tilde{S}\tilde{W}^{\mathrm{t}}$ are two sign component decompositions (2.3) with the same inner dimension r. We say that the decompositions are *equivalent* if there is a signed permutation matrix Π for which $\tilde{S} = S\Pi$ and $\tilde{W} = W\Pi$.

Alternatively, consider two sign component decompositions $B = \sum_{i=1}^{r} s_i w_i^{\mathrm{t}}$ and $B = \sum_{i=1}^{r} \tilde{s}_i \tilde{w}_i$ with the same number r of terms. The decompositions are *equivalent* if there is a signed permutation $(\pi, \boldsymbol{\xi})$ on r letters for which $\tilde{s}_i = \xi_i s_{\pi(i)}$ and $\tilde{w}_i = \xi_i w_{\pi(i)}$ for each $i = 1, \dots, r$.

4.4. **The role of Schur independence.** As we have just seen, signed permutations preserve the class of sign component decompositions of a given matrix. Meanwhile, the proof of Proposition 4.1 warns us that we can sometimes map one sign component decomposition to an inequivalent decomposition via an invertible transformation. Remarkably, we can preclude the latter phenomenon by narrowing our attention to Schur independent sign matrices. In this case, sign permutations are the *only* invertible transformations that respect the sign structure.

Proposition 4.3 (Schur independence: Transformations)**.** *Let $S \in \{\pm 1\}^{n \times r}$ be a Schur independent sign matrix, and let $Q \in \mathbb{R}^{r \times r}$ be an invertible matrix. Then $SQ \in \{\pm 1\}^{n \times r}$ is a sign matrix if and only if Q is a signed permutation.*

Proof. If Q is a signed permutation, then it is immediate that SQ is a sign matrix. The reverse implication is the more interesting fact.

Introduce notation for the columns of the matrices under discussion:

$$S = \begin{bmatrix} s_1 & \dots & s_r \end{bmatrix} \in \{\pm 1\}^{n \times r} \quad \text{and} \quad Q = \begin{bmatrix} q_1 & \dots & q_r \end{bmatrix} \in \mathbb{R}^{r \times r} \quad \text{and} \quad SQ = \begin{bmatrix} \tilde{s}_1 & \dots & \tilde{s}_r \end{bmatrix} \in \{\pm 1\}^{n \times r}.$$

For each index $1 \le k \le r$, the kth column $\tilde{s}_k$ of the matrix SQ satisfies

$$\tilde{s}_k = Sq_k = \sum_{i=1}^{r} \langle \mathbf{e}_i, \ q_k \rangle s_i.$$

By assumption, $\tilde{\boldsymbol{s}}_k$ is a sign vector, so

$$\boldsymbol{e} = \tilde{\boldsymbol{s}}_k \odot \tilde{\boldsymbol{s}}_k = \sum\nolimits_{i,j=1}^r \langle \boldsymbol{e}_i,\ \boldsymbol{q}_k \rangle \langle \boldsymbol{e}_j,\ \boldsymbol{q}_k \rangle (\boldsymbol{s}_i \odot \boldsymbol{s}_j) = \left(\sum\nolimits_{i=1}^r \langle \boldsymbol{e}_i,\ \boldsymbol{q}_k \rangle^2\right) \boldsymbol{e} + 2 \sum\nolimits_{i<j} \langle \boldsymbol{e}_i,\ \boldsymbol{q}_k \rangle \langle \boldsymbol{e}_j,\ \boldsymbol{q}_k \rangle (\boldsymbol{s}_i \odot \boldsymbol{s}_j).$$

Schur independence of the matrix $\boldsymbol{S}$ ensures that the family $\{\boldsymbol{e}\} \cup \{\boldsymbol{s}_i \odot \boldsymbol{s}_j : i < j\} \subset \mathbb{R}^r$ is linearly independent. As a consequence,

$$\sum\nolimits_{i=1}^r \langle \boldsymbol{e}_i,\ \boldsymbol{q}_k \rangle^2 = 1 \quad \text{and} \quad \langle \boldsymbol{e}_i,\ \boldsymbol{q}_k \rangle \langle \boldsymbol{e}_j,\ \boldsymbol{q}_k \rangle = 0 \quad \text{when } i \neq j.$$

Since $\boldsymbol{q}_k$ solves this quadratic system, it must be a signed standard basis vector: $\boldsymbol{q}_k = \xi_k \boldsymbol{e}_{\pi(k)} \in \mathbb{R}^r$ for a sign $\xi_k \in \{\pm 1\}$ and an index $\pi(k) \in \{1,\dots,r\}$. Since the matrix $\boldsymbol{Q}$ is invertible, it must be the case that π is a permutation on r letters. It follows that $\boldsymbol{Q}$ is a signed permutation. $\qquad\square$

4.5. **Uniqueness.** With this preparation, we can delineate circumstances where the (minimal) sign component decomposition of a low-rank matrix is unique up to equivalence.

Theorem 4.4 (Sign component decomposition: Uniqueness). *Consider a matrix $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ that admits a sign component decomposition $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{W}^{\mathrm{t}}$. Assume that*

   (1) The sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ is Schur independent;
   (2) The weight matrix $\boldsymbol{W} \in \mathbb{R}^{m \times r}$ has full column rank.

Then all minimal sign component decompositions of $\boldsymbol{B}$ (with inner dimension r) are equivalent.

Proof. The sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ has full column rank because it is Schur independent (Fact 2.3(1)), while the weight matrix $\boldsymbol{W} \in \mathbb{R}^{m \times r}$ has full column rank by assumption. We discover that the matrix $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{W}^{\mathrm{t}}$ has rank r. Therefore, every sign component decomposition of $\boldsymbol{B}$ has inner dimension at least r, and the distinguished decomposition has the minimal inner dimension.

Suppose that $\boldsymbol{B} = \tilde{\boldsymbol{S}}\tilde{\boldsymbol{W}}^{\mathrm{t}}$ is another sign component decomposition with inner dimension r. Since $\boldsymbol{B}$ has rank r, both factors $\tilde{\boldsymbol{S}}$ and $\tilde{\boldsymbol{W}}$ must have full column rank. As a consequence, there is an invertible transformation $\boldsymbol{Q} \in \mathbb{R}^{r \times r}$ for which $\tilde{\boldsymbol{S}} = \boldsymbol{S}\boldsymbol{Q}$. Since $\boldsymbol{S}$ is a Schur independent sign matrix and $\tilde{\boldsymbol{S}}$ is a sign matrix, Proposition 4.3 forces $\boldsymbol{Q}$ to be a signed permutation. Now, we have the chain of identities

$$\boldsymbol{S}\boldsymbol{W}^{\mathrm{t}} = \boldsymbol{B} = \tilde{\boldsymbol{S}}\tilde{\boldsymbol{W}}^{\mathrm{t}} = \boldsymbol{S}\boldsymbol{Q}\tilde{\boldsymbol{W}}^{\mathrm{t}} = \boldsymbol{S}(\tilde{\boldsymbol{W}}\boldsymbol{Q}^{\mathrm{t}})^{\mathrm{t}}.$$

Since the matrix $\boldsymbol{S}$ has full column rank, we can cancel $\boldsymbol{S}$ to see that $\tilde{\boldsymbol{W}}\boldsymbol{Q}^{\mathrm{t}} = \boldsymbol{W}$. The signed permutation $\boldsymbol{Q}$ is orthogonal, so it follows that $\tilde{\boldsymbol{W}} = \boldsymbol{W}\boldsymbol{Q}$.

To summarize, we have been given two sign component decompositions $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{W}^{\mathrm{t}} = \tilde{\boldsymbol{S}}\tilde{\boldsymbol{W}}^{\mathrm{t}}$ with inner dimension r. We have shown that they are related by $\tilde{\boldsymbol{S}} = \boldsymbol{S}\boldsymbol{Q}$ and $\tilde{\boldsymbol{W}} = \boldsymbol{W}\boldsymbol{Q}$ for a signed permutation $\boldsymbol{Q}$. Therefore, the two decompositions are equivalent. $\qquad\square$

Theorem 4.4 describes conditions under which the minimal sign component decomposition of a matrix is uniquely determined. It is natural to demand that both the left and the right factors have full column rank. The geometry of the factorization problem dictates the stronger requirement that the sign matrix is Schur independent. As we have discussed, most families of $r < \frac{1}{2}(1 + \sqrt{8n-7})$ sign vectors are Schur independent, so this condition holds for a rich class of matrices.

5. Computation of the asymmetric sign component decomposition

In this section, we derive and justify Algorithm 3, which computes the asymmetric sign component decomposition of a matrix whose sign component is Schur independent. We establish the following result.

Theorem 5.1 (Sign component decomposition: Computation). *Consider a matrix $\boldsymbol{B} \in \mathbb{R}^{n \times m}$ that admits a sign component decomposition $\boldsymbol{B} = \boldsymbol{S}\boldsymbol{W}^{\mathrm{t}}$. Assume that*

   (1) The sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ is Schur independent;
   (2) The weight matrix $\boldsymbol{W} \in \mathbb{R}^{m \times r}$ has full column rank.

Then, with probability one, Algorithm 1 identifies the minimal sign component decomposition, up to signed permutation. That is, the output is a pair $(\tilde{S}, \tilde{W}) \in \{\pm 1\}^{n \times r} \times \mathbb{R}^{m \times r}$ where $\tilde{S} = S\Pi$ and $\tilde{W} = W\Pi$ for a signed permutation $\Pi \in \mathbb{R}^{r \times r}$.

We prove Theorem 5.1 below in Section 5.2.

5.1. **Factorization and semidefinite programming.**

Although constrained matrix factorization is viewed as a challenging problem, certain aspects are simpler than they appear. In particular, we can expose properties of the components of a matrix factorization by means of a semidefinite constraint.

Fact 5.2 (Factorization constraint). *Let $B \in \mathbb{R}^{n \times m}$ be a matrix. The semidefinite relation*

$$\begin{bmatrix} X & B \\ B^{\mathrm{t}} & Y \end{bmatrix} \succcurlyeq \mathbf{0} \tag{5.1}$$

enforces a factorization of B in the following sense.

(1) If $B = UV^{\mathrm{t}}$, then (5.1) holds when $X = UU^{\mathrm{t}}$ and $Y = VV^{\mathrm{t}}$.

(2) If (5.1) holds, then we can decompose $B = UV^{\mathrm{t}}$ into factors $U \in \mathbb{R}^{n \times r}$ and $V \in \mathbb{R}^{m \times r}$ that satisfy $X = UU^{\mathrm{t}}$ and $Y = VV^{\mathrm{t}}$. The inner dimension meets the bound $r \leq m + n$.

We omit the easy proof, because we do not use this result directly.

The factorization constraint (5.1) does not give us direct access to the factors U and V. Nevertheless, we can place restrictions on the variables X and Y to limit the possible values that the factors U and V can take. If the conditions are strong enough, it is sometimes possible to determine the factors completely, modulo symmetries.

Example 5.3 (From SVD to eigenvalue decomposition). Let $B \in \mathbb{R}^{n \times m}$ be a matrix. Consider the semidefinite program

$$\underset{X \in \mathbb{H}_n, Y \in \mathbb{H}_m}{\text{minimize}} \quad \tfrac{1}{2}(\operatorname{trace}(X) + \operatorname{trace}(Y)) \quad \text{subject to} \quad \begin{bmatrix} X & B \\ B^{\mathrm{t}} & Y \end{bmatrix} \succcurlyeq \mathbf{0}.$$

Every minimizer takes the form $X_\star = U\Sigma U^{\mathrm{t}}$ and $Y_\star = V\Sigma V^{\mathrm{t}}$ where $B = U\Sigma V^{\mathrm{t}}$ is a singular value decomposition. We can find the left and right singular vectors (U, V) of B by computing the eigenvalue decompositions of $X_\star$ and $Y_\star$. As a side note, the minimal value of the optimization problem is the Schatten 1-norm (i.e., the sum of singular values) of the matrix B.

As we will see, a more elaborate version of the procedure in Example 5.3 allows us to compute an asymmetric sign component decomposition. To develop this approach, we require ingredients (Fact 3.3 and Fact 3.4) from our work on symmetric sign component decomposition.

5.2. **Overview of algorithm and proof of Theorem 5.1.**

Given an input matrix $B = SW^{\mathrm{t}}$ with a Schur independent sign component S, our aim is to find the (unknown) asymmetric sign component decomposition. We reduce this challenge to the solved problem of computing a symmetric sign component decomposition of a correlation matrix. In this section, we outline the procedure, along with the proof of Theorem 5.1. Algorithm 1 encapsulates the computations, and some details of the argument are postponed to the next sections.

The first step is to construct a correlation matrix whose symmetric sign component decomposition has the same sign factor as the input matrix B. To that end, construct the orthogonal projector $P \in \mathbb{H}_n$ onto the range of B. Then solve the semidefinite program (SDP)

$$\begin{aligned}
\underset{X \in \mathbb{H}_n, Y \in \mathbb{H}_m}{\text{minimize}} \quad & \operatorname{trace}(Y) \\
\text{subject to} \quad & \operatorname{trace}(PX) = n \text{ and } \operatorname{diag}(X) = \mathbf{e}; \\
& \begin{bmatrix} X & B \\ B^{\mathrm{t}} & Y \end{bmatrix} \succcurlyeq \mathbf{0}.
\end{aligned} \tag{5.2}$$

Fact 5.2 shows that the semidefinite constraint in (5.2) links the variables X and Y to a factorization of B. Meanwhile, courtesy of Fact 3.4, the equality constraints in (5.2) force the variable X to be a correlation matrix whose range equals the range of B. The following lemma packages these claims.

Proposition 5.4 (Factorization SDP)**.** *Instate the assumptions of Theorem 5.1. Let $(X_\star, Y_\star)$ be the unique minimizer of the optimization problem* (5.2). *Then $X_\star = S \operatorname{diag}(\tau) S^{\mathrm{t}}$ where $S \in \{\pm 1\}^{n \times r}$ is the sign component of B and $\tau \in \Delta_r^+$.*

We prove a more detailed version of Proposition 5.4 below in Section 5.4.

The next step is to extract the sign component of the correlation matrix $X_\star$ that solves (5.2). According to Proposition 5.4, the correlation matrix $X_\star$ meets the requirements of Fact 3.3. Therefore, we can invoke Algorithm 3, the symmetric sign component decomposition method, to obtain a factorization

$$X_\star = \tilde{S} \operatorname{diag}(\tilde{\tau}) \tilde{S}^{\mathrm{t}} \quad \text{where} \quad \tilde{S} = S\Pi \text{ for a signed permutation } \Pi.$$

We cannot resolve the signed permutation, but the computed sign component $\tilde{S}$ is equivalent with the designated sign component S.

To complete the sign component decomposition, it remains to determine the weight matrix. We may do so by solving the linear system

$$\text{find } \tilde{W} \in \mathbb{R}^{m \times r} \quad \text{subject to} \quad B = \tilde{S}\tilde{W}^{\mathrm{t}}.$$

The solution exists because $B = SW^{\mathrm{t}} = \tilde{S}(W\Pi)^{\mathrm{t}}$. The solution $\tilde{W} = W\Pi$ is unique because $\tilde{S}$ has Schur independent columns, and so its columns are also linearly independent (Fact 2.3(1)).

The pair $(\tilde{S}, \tilde{W})$ yields a sign component decomposition of the matrix B that is equivalent with the specified decomposition $B = SW^{\mathrm{t}}$. This observation completes the proof of Theorem 5.1.

5.3. **Positive-semidefinite matrices.**

It remains to establish Proposition 5.4. The argument depends on core properties of psd matrices, which we collect here. For references, see [Bha97, Bha07].

Fact 5.5 (Conjugation rule)**.** *Conjugation respects the semidefinite order in the following sense.*

  (1) If $X \succcurlyeq 0$, then $KXK^{\mathrm{t}} \succcurlyeq 0$ for each matrix K with compatible dimensions.
  (2) If K has full column rank and $KXK^{\mathrm{t}} \succcurlyeq 0$, then $X \succcurlyeq 0$.

Fact 5.6 (Schur complements)**.** *Assume that $X \in \mathbb{H}_n$ is a (strictly) positive-definite matrix. Then*

$$\begin{bmatrix} X & K \\ K^{\mathrm{t}} & Y \end{bmatrix} \succcurlyeq 0 \quad \text{if and only if} \quad Y \succcurlyeq K^{\mathrm{t}} X^{-1} K.$$

Related results hold when X is merely psd.

Fact 5.7 (Trace is monotone)**.** *Let X and Y be psd matrices that satisfy $X \succcurlyeq Y$. Then $\operatorname{trace}(X) \geq \operatorname{trace}(Y)$, and equality holds precisely when $X = Y$.*

5.4. **The Factorization SDP.**

We are now prepared to prove Proposition 5.4, which describes the solution of the factorization SDP (5.2). The proposition follows instantly from a more precise lemma.

Lemma 5.8 (Factorization SDP)**.** *Instate the assumptions of Theorem 5.1. Construct the orthogonal projector $P \in \mathbb{H}_n$ onto the range of B. Define the positive-definite diagonal matrix*

$$D = \operatorname{diag}(\|w_1\|_{\ell_2}, \dots, \|w_r\|_{\ell_2}) \succ 0 \quad \text{where} \quad W = \begin{bmatrix} w_1 & \dots & w_r \end{bmatrix}.$$

Then the unique solution to the semidefinite optimization problem (5.2) *is the pair*

$$X_\star = (\operatorname{trace} D)^{-1} SDS^{\mathrm{t}} \quad \text{and} \quad Y_\star = (\operatorname{trace} D) WD^{-1}W^{\mathrm{t}}.$$

Proof. Recall that $B = SW^{\mathrm{t}}$ for a Schur independent sign matrix S and a matrix W with full column rank.

First, we argue that a feasible point X of the factorization SDP (5.2) must be a correlation matrix of the form

$$X = S \operatorname{diag}(\tau) S^{\mathrm{t}} \quad \text{for } \tau \in \Delta_r. \tag{5.3}$$

Indeed, the block matrix constraint in (5.2) ensures that $X \succcurlyeq \mathbf{0}$, and the constraint $\mathrm{diag}(X) = \mathbf{e}$ makes X a correlation matrix. At the same time, since the matrix $W \in \mathbb{R}^{m \times r}$ has full column rank,

$$\mathrm{range}(P) = \mathrm{range}(B) = \mathrm{range}(SW^{\mathrm{t}}) = \mathrm{range}(S).$$

Fact 3.4 shows that the constraint $\mathrm{trace}(PX) = n$ isolates the family $\{S \, \mathrm{diag}(\boldsymbol{\tau}) \, S^{\mathrm{t}} : \boldsymbol{\tau} \in \Delta_r\}$ of correlation matrices. This establishes the claim.

Next, substitute the expression (5.3) into the block matrix constraint in (5.2) and use the condition $B = SW^{\mathrm{t}}$ to factorize:

$$\begin{bmatrix} X & B \\ B^{\mathrm{t}} & Y \end{bmatrix} = \begin{bmatrix} S & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathrm{diag}(\boldsymbol{\tau}) & W^{\mathrm{t}} \\ W & Y \end{bmatrix} \begin{bmatrix} S & \mathbf{0} \\ \mathbf{0} & \mathbf{I} \end{bmatrix}^{\mathrm{t}} \succcurlyeq \mathbf{0}.$$

Since S is Schur independent, it has full column rank (Fact 2.3(1)). Therefore, the conjugation rule (Fact 5.5) implies that the psd constraint in the last display is equivalent with the condition

$$\begin{bmatrix} \mathrm{diag}(\boldsymbol{\tau}) & W^{\mathrm{t}} \\ W & Y \end{bmatrix} \succcurlyeq \mathbf{0}. \tag{5.4}$$

Now, we can recognize that $\mathrm{diag}(\boldsymbol{\tau})$ is a strictly positive-definite matrix. Indeed, owing to (5.4), the relation $\tau_i = 0$ would imply that the corresponding column $\boldsymbol{w}_i$ of the weight matrix equals zero, but this is impossible because W has full column rank.

Apply the Schur complement rule (Fact 5.6) to the matrix (5.4) to confirm that

$$Y \succcurlyeq W \, \mathrm{diag}(\boldsymbol{\tau})^{-1} W^{\mathrm{t}} \succ \mathbf{0}.$$

The objective function, $\mathrm{trace}(Y)$, of the semidefinite program (5.2) is strictly monotone with respect to the semidefinite order (Fact 5.7). The variable Y is otherwise unconstrained, so the SDP achieves its minimum if and only if

$$Y = W \, \mathrm{diag}(\boldsymbol{\tau})^{-1} W^{\mathrm{t}}.$$

It remains to determine the vector $\boldsymbol{\tau}_\star \in \Delta_r^+$ that minimizes the trace of Y.

To that end, calculate that

$$\mathrm{trace}(Y) = \mathrm{trace}(W \, \mathrm{diag}(\boldsymbol{\tau})^{-1} W^{\mathrm{t}}) = \sum_{i=1}^r \tau_i^{-1} \|\boldsymbol{w}_i\|_{\ell_2}^2 \geq \min_{1 \leq k \leq r} \left(\tau_k^{-1} \|\boldsymbol{w}_k\|_{\ell_2} \right) \cdot \sum_{i=1}^r \|\boldsymbol{w}_i\|_{\ell_2} > 0.$$

Equality holds if and only if the quantities $\tau_k^{-1} \|\boldsymbol{w}_k\|_{\ell_2}$ are identical for all indices k. Since $\boldsymbol{\tau} \in \Delta_r^+$, we may conclude that the minimizer $\boldsymbol{\tau}_\star$ has coordinates

$$(\boldsymbol{\tau}_\star)_k = \|\boldsymbol{w}_k\|_{\ell_2} \left(\sum_{i=1}^r \|\boldsymbol{w}_i\|_{\ell_2} \right)^{-1} \quad \text{for each index } k.$$

In summary, we have shown that the unique matrices that optimize (5.2) take the form

$$X_\star = S \, \mathrm{diag}(\boldsymbol{\tau}_\star) \, S^{\mathrm{t}} \quad \text{and} \quad Y_\star = W \, \mathrm{diag}(\boldsymbol{\tau}_\star)^{-1} W^{\mathrm{t}}.$$

Identify the diagonal matrix D from the statement to complete the proof. $\qquad \square$

6. ASYMMETRIC BINARY COMPONENT DECOMPOSITION

In this section, we develop a procedure (Algorithm 2) for computing an asymmetric binary component decomposition (2.5)–(2.6). We prove Theorem II, which states that the algorithm succeeds under a Schur independence condition. Our approach reduces the problem of computing a binary component decomposition to the problem of computing a sign component decomposition of a related matrix.

6.1. **Correspondence between binary vectors and sign vectors.** As we have discussed, there is a one-to-one correspondence F between sign vectors and binary vectors (2.7). The correspondence between asymmetric sign component decompositions and binary component decompositions, however, is more subtle because they are invariant under different transformation. Indeed, $\boldsymbol{s}\boldsymbol{w}^{\mathrm{t}}$ does not change if we flip the sign of both $\boldsymbol{s} \in \{\pm 1\}^n$ and $\boldsymbol{w} \in \mathbb{R}^m$. On the other hand, the matrix $\boldsymbol{z}\boldsymbol{w}^{\mathrm{t}}$ completely determines the vectors $\boldsymbol{z} \in \{0,1\}^n$ and $\boldsymbol{w} \in \mathbb{R}^m$.

6.2. **Reducing binary component decomposition to sign component decomposition.** Given a matrix that has a binary component decomposition, we can apply a simple transformation to construct a related matrix that admits a sign component decomposition

Proposition 6.1 (Binary component decomposition: Reduction). *Consider a matrix $C \in \mathbb{R}^{n \times m}$ that has a binary component decomposition*

$$C = ZW^{\mathsf{t}} \quad \text{where} \quad Z \in \{0,1\}^{n \times r} \text{ and } W \in \mathbb{R}^{m \times r}.$$

Construct the matrix

$$B = F(C) = 2C - \mathbf{E} \in \mathbb{R}^{n \times m}.$$

Then B admits a sign component decomposition with inner dimension $r + 1$:

$$B = \begin{bmatrix} S & \mathbf{e} \end{bmatrix} \begin{bmatrix} W & W\mathbf{e} - \mathbf{e} \end{bmatrix}^{\mathsf{t}} \quad \text{where} \quad S = F(Z). \tag{6.1}$$

Recall that $\mathbf{E} = \mathbf{e}\mathbf{e}^{\mathsf{t}}$ is a matrix of ones with appropriate dimensions.

Proof. The result follows from a straightforward calculation:

$$B = 2C - \mathbf{E} = 2ZW^{\mathsf{t}} - \mathbf{E} = (2Z - \mathbf{E})W^{\mathsf{t}} + \mathbf{E}W^{\mathsf{t}} - \mathbf{E}$$

$$= F(Z)W^{\mathsf{t}} + \mathbf{e}(W\mathbf{e} - \mathbf{e})^{\mathsf{t}} = \begin{bmatrix} F(Z) & \mathbf{e} \end{bmatrix} \begin{bmatrix} W & W\mathbf{e} - \mathbf{e} \end{bmatrix}^{\mathsf{t}}.$$

Recognize the matrix $S = F(Z)$ to complete the argument. $\square$

6.3. **Resolving the sign ambiguity.** Proposition 6.1 allows us to reduce the problem of computing a binary component decomposition to the problem of computing a sign component decomposition. Nevertheless, the sign component decomposition has a sign invariance that is not present in the binary component decomposition. The next result explains how to resolve this ambiguity.

Proposition 6.2 (Sign ambiguity). *Instate the notation of Proposition 6.1. Assume that the unique sign component decomposition of B with inner dimension $r + 1$ is the one given in (6.1), which we write as*

$$B = \begin{bmatrix} s_1 & \dots & s_r & \mathbf{e} \end{bmatrix} \begin{bmatrix} w_1 & \dots & w_r & W\mathbf{e} - \mathbf{e} \end{bmatrix}^{\mathsf{t}}. \tag{6.2}$$

Suppose that we have computed another sign component decomposition

$$B = \tilde{S}\tilde{W}^{\mathsf{t}} = \begin{bmatrix} \tilde{s}_1 & \dots & \tilde{s}_r & \mathbf{e} \end{bmatrix} \begin{bmatrix} \tilde{w}_1 & \dots & \tilde{w}_r & \tilde{w}_{r+1} \end{bmatrix}^{\mathsf{t}}.$$

By uniqueness, there is a signed permutation $(\pi, \boldsymbol{\xi})$ on r letters for which

$$\tilde{s}_i = \xi_i s_{\pi(i)} \quad \text{and} \quad \tilde{w}_i = \xi_i w_{\pi(i)} \quad \text{for each index } i = 1, \dots, r.$$

Then the sign vector $\boldsymbol{\xi} \in \mathbb{R}^r$ is the unique solution to the linear system

$$\begin{bmatrix} \tilde{w}_1 & \dots & \tilde{w}_r \end{bmatrix} \boldsymbol{\xi} = \tilde{w}_{r+1} + \mathbf{e}$$

Moreover, the binary matrix satisfies

$$Z = \begin{bmatrix} z_1 & \dots & z_r \end{bmatrix} \quad \text{where} \quad z_{\pi(i)} = F^{-1}(\xi_i \tilde{s}_i).$$

Proof. If the sign component decomposition (6.2) with $r + 1$ terms is determined up to signed permutation, then $\pm\mathbf{e}$ must appear among the columns of $\tilde{S}$. By sign change and permutation, we may arrange that the last column of $\tilde{S}$ equals $\mathbf{e}$. The simultaneous transformation on $\tilde{W}$ ensures that its last column satisfies $\tilde{w}_{r+1} = W\mathbf{e} - \mathbf{e}$. The uniqueness of the decomposition (6.2) ensures that the weight matrix has full column rank, and so $\tilde{W}$ also has full column rank. Now, observe that

$$\begin{bmatrix} \tilde{w}_1 & \dots & \tilde{w}_r \end{bmatrix} \boldsymbol{\xi} = \begin{bmatrix} \xi_1 w_{\pi(1)} & \dots & \xi_r w_{\pi(r)} \end{bmatrix} \boldsymbol{\xi} = \sum_{i=1}^r \xi_i^2 w_{\pi(i)} = \sum_{i=1}^r w_i = W\mathbf{e} = \tilde{w}_{r+1} + \mathbf{e}.$$

The solution to this linear system is uniquely determined, so it must equal the true vector $\boldsymbol{\xi}$ of sign flips. Therefore,

$$z_{\pi(i)} = F^{-1}(s_{\pi(i)}) = F^{-1}(\xi_i \tilde{s}_i) \quad \text{for each index } i = 1, \dots, r.$$

This completes the argument. $\square$

6.4. **Uniqueness.** Next, we confirm that binary component decompositions are unique under a Schur independence condition. The key step is to argue that ordinary permutations are the only invertible transformations that preserve Schur independence of binary vectors. This result is an analog of Proposition 4.3, and the proof follows a similar pattern.

Proposition 6.3 (Binary Schur independence: Transformations)**.** *Let $Z \in \{0, 1\}^{n \times r}$ be a Schur independent binary matrix, and let $Q \in \mathbb{R}^{r \times r}$ be an invertible matrix. Then ZQ is a binary matrix if and only if Q is a permutation matrix.*

Proof. If Q is a permutation, then it is clear that ZQ is a binary matrix. Let us prove the converse.

Write z_i for the columns of Z; write q_i for the columns of Q; and write $\tilde{z}_i$ for the columns of ZQ. For each index $1 \le k \le r$,

$$\tilde{z}_k = Zq_k = \sum_{i=1}^{r} \langle \mathbf{e}_i, \, q_k \rangle z_i.$$

The vector $\tilde{z}_k$ is binary, so

$$\mathbf{0} = (\mathbf{e} - \tilde{z}_k) \odot \tilde{z}_k = \sum_{i=1}^{r} \langle \mathbf{e}_i, \, q_k \rangle (\mathbf{e} \odot z_i) - \sum_{i,j=1}^{r} \langle \mathbf{e}_i, \, q_k \rangle \langle \mathbf{e}_j, \, q_k \rangle (z_i \odot z_j)$$

$$= \sum_{i=1}^{r} \langle \mathbf{e}_i, \, q_k \rangle (1 - \langle \mathbf{e}_i, \, q_k \rangle) z_i - 2 \sum_{i<j} \langle \mathbf{e}_i, \, q_k \rangle \langle \mathbf{e}_j, \, q_k \rangle (z_i \odot z_j).$$

We have used the fact that $z_i = \mathbf{e} \odot z_i = z_i \odot z_i$ for each binary vector. Schur independence of the matrix Z ensures that the vectors on the right-hand side of this expression compose a linearly independent family. It follows that

$$\langle \mathbf{e}_i, \, q_k \rangle (1 - \langle \mathbf{e}_i, \, q_k \rangle) = 0 \quad \text{and} \quad \langle \mathbf{e}_i, \, q_k \rangle \langle \mathbf{e}_j, \, q_k \rangle = 0 \quad \text{when } i \ne j.$$

Therefore, q_k must be a standard basis vector: $q_k = \mathbf{e}_{\pi(k)} \in \mathbb{R}^r$ for an index $\pi(k) \in \{1, \dots, r\}$. Since the matrix Q is invertible, π is a permutation on r letters. In other words, Q is a permutation matrix. $\qquad\square$

As a consequence, we obtain a result about the uniqueness of asymmetric binary component decompositions, modulo permutation.

Theorem 6.4 (Binary component decomposition: Uniqueness)**.** *Consider a matrix $C \in \mathbb{R}^{n \times m}$ that admits a binary component decomposition $C = ZW^{\mathrm{t}}$. Assume that*

(1) The binary matrix $Z \in \{0, 1\}^{n \times r}$ is Schur independent;
(2) The weight matrix $W \in \mathbb{R}^{m \times r}$ has full column rank.

Then all minimal sign component decompositions of C (with inner dimension r) are equivalent, up to simultaneous permutation of the columns of the factors.

We omit the proof, which mirrors the argument in Theorem 6.4.

6.5. **Computation.** Proposition 6.1 and Proposition 6.2 give us a mechanism for computing a binary component decomposition, provided that an associated matrix has a unique sign component decomposition. We can exploit our theory on the tractable computation of sign component decompositions to identify situations where we can compute binary component decompositions.

Theorem 6.5 (Binary component decomposition: Computation)**.** *Consider a matrix $C \in \mathbb{R}^{n \times m}$ that admits a sign component decomposition $C = ZW^{\mathrm{t}}$. Assume that*

(1) The sign matrix $Z \in \{0, 1\}^{n \times r}$ is Schur independent;
(2) The weight matrix $W \in \mathbb{R}^{m \times r}$ has full column rank.

Then, with probability one, Algorithm 2 identifies the minimal binary component decomposition.

Proof of Theorem 6.5. Suppose that $C \in \mathbb{R}^{n \times m}$ has a binary component decomposition $C = \sum_{i=1}^{r} z_i w_i^{\mathrm{t}}$ involving a Schur independent family $\{z_1, \dots, z_r\} \subseteq \{0, 1\}^r$ of binary components. Introduce the associated sign vectors $s_i = F(z_i)$. By Fact 2.5, the family $\{s_1, \dots, s_r, \mathbf{e}\}$ of sign vectors is Schur independent.

Proposition 6.1 shows that the matrix $B = 2C - \mathbf{E} = \begin{bmatrix} S & \mathbf{e} \end{bmatrix} \begin{bmatrix} W & W\mathbf{e} - \mathbf{e} \end{bmatrix}^{\mathrm{t}}$. Define $s_{r+1} = \mathbf{e}$ and $w_{r+1} = W\mathbf{e} - \mathbf{e}$. By Theorem 5.1, Algorithm 1 allows us to compute pairs $(\tilde{s}_i, \tilde{w}_i)$ with the property that $\tilde{s}_i = \xi_i s_{\pi(i)}$ and $\tilde{w}_i = \xi_i w_{\pi(i)}$ for each $i = 1, \dots, r+1$ where (π, ξ) is a signed permutation. By change of sign and

permutation, we can assume that $\tilde{\boldsymbol{s}}_{r+1} = \mathbf{e}$ and $\tilde{\boldsymbol{w}}_{r+1} = \boldsymbol{W}\mathbf{e} - \mathbf{e}$. Proposition 6.2 shows that we can use these computed weight vectors $\tilde{\boldsymbol{w}}_i$ to find the sign vector $\boldsymbol{\xi}$, and we obtain the binary components of $\boldsymbol{C}$ as $\boldsymbol{z}_{\pi(i)} = \tilde{\boldsymbol{z}}_i = \boldsymbol{F}^{-1}(\xi_i \tilde{\boldsymbol{s}}_i)$.

Last, we define the binary factor matrix $\tilde{\boldsymbol{Z}} = \begin{bmatrix} \tilde{\boldsymbol{z}}_1 & \dots & \tilde{\boldsymbol{z}}_r \end{bmatrix}$ and construct a weight matrix with the correct signs: $\tilde{\boldsymbol{W}}_+ = \begin{bmatrix} \xi_1 \tilde{\boldsymbol{w}}_1 & \dots & \xi_r \tilde{\boldsymbol{w}}_r \end{bmatrix}$. □

7. ROBUSTNESS

In this paper, we have designed factorization algorithms to operate on low-rank matrices that admit an exact sign component decomposition or binary component decomposition. In most practical applications, however, the data matrix will be contaminated by noise, errors, or outliers. One way to handle these non-idealities is to process the data to remove corruptions. Afterward, we can apply our algorithms to factorize the clean data matrix. This section describes some situations where we can implement this idea.

7.1. **Approximate factorization of a noisy matrix.** We focus on computing an asymmetric sign component decomposition of a noisy data matrix:
$$\boldsymbol{B} = \boldsymbol{L}_0 + \boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m} \quad \text{where} \quad \boldsymbol{L}_0 = \boldsymbol{S}\boldsymbol{W}^{\mathsf{t}}.$$
As usual, the factor $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ and $\boldsymbol{W} \in \mathbb{R}^{m \times r}$. The matrix $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m}$ captures the noise.

Given the matrix $\boldsymbol{B}$, our computational goal is to remove the noise $\boldsymbol{\Omega}_0$ completely. Then we can factorize the clean matrix $\boldsymbol{L}_0$ that remains. In summary,

(1) **Denoise:** Remove the noise $\boldsymbol{\Omega}_0$ from the observed matrix $\boldsymbol{B}$ to obtain the clean data $\boldsymbol{L}_0$.
(2) **Factorize:** Compute a sign component decomposition $\boldsymbol{L}_0 = \tilde{\boldsymbol{S}}\tilde{\boldsymbol{W}}^{\mathsf{t}}$.

We can realize this strategy for several different noise models. In particular,

- **Gross errors in entries:** The corruption $\boldsymbol{\Omega}_0$ is a sparse matrix. This model is appropriate for errors in individual measurements.
- **Gross errors in columns:** The corruption $\boldsymbol{\Omega}_0$ is a column-sparse matrix. This model is appropriate for handling outliers.

To facilitate the analysis, we also instate simple probabilistic models for the clean data matrix $\boldsymbol{L}_0$ and for the noise $\boldsymbol{\Omega}_0$. These assumptions can be relaxed substantially.

7.2. **The Gaussian loadings model.** We work with a generative probabilistic model for the clean data matrix. This model combines a fixed set of sign vectors with random coefficients.

Model 7.1 (Gaussian loadings)**.** Select a Schur-independent sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ and a natural number $m \geq r$. The Gaussian loading model GLM($\boldsymbol{S}, m$) is a distribution on matrices in $\mathbb{R}^{n \times m}$. A matrix $\boldsymbol{L}_0$ drawn from this model takes the form
$$\boldsymbol{L}_0 = \frac{1}{\sqrt{r}}\boldsymbol{S}\boldsymbol{G}^{\mathsf{t}} \quad \text{where} \quad \boldsymbol{G} \in \mathbb{R}^{m \times r} \text{ is standard normal.}$$

Owing to Theorem I, a realization $\boldsymbol{L}_0 \sim \text{GLM}(\boldsymbol{S}, m)$ almost surely has a unique sign component decomposition. Moreover, we can produce an equivalent decomposition $\boldsymbol{L}_0 = \tilde{\boldsymbol{S}}\tilde{\boldsymbol{W}}^{\mathsf{t}}$ using Algorithm 1.

Suppose we are given a noisy observation of a matrix drawn from the Gaussian loadings model:
$$\boldsymbol{B} = \boldsymbol{L}_0 + \boldsymbol{\Omega}_0 \quad \text{where} \quad \boldsymbol{L}_0 \sim \text{GLM}(\boldsymbol{S}, m).$$
We are interested in removing the noise. The difficulty of this problem depends on the choice of the sign matrix $\boldsymbol{S}$, the number m of independent data points that we sample, and the type of noise $\boldsymbol{\Omega}_0$ that we must contend with.

Let us explain how the geometry of the Gaussian loadings model enters the picture. The sign matrix $\boldsymbol{S}$ determines the anisotropy of the columns of $\boldsymbol{L}_0$. Indeed,
$$\mathbb{E}[\boldsymbol{L}_0\boldsymbol{L}_0^{\mathsf{t}}] = \frac{1}{r}\mathbb{E}[\boldsymbol{S}\boldsymbol{G}\boldsymbol{G}^{\mathsf{t}}\boldsymbol{S}^{\mathsf{t}}] = \frac{m}{r} \cdot \boldsymbol{S}\boldsymbol{S}^{\mathsf{t}} \in \mathbb{H}_n.$$

We can summarize how well the columns of L_0 fill out their span by means of the *permeance statistic*:

$$\nu(S) = \frac{1}{n}\lambda_{\min}(S^{\mathrm{t}}S) = \frac{1}{n}\min_{\|u\|_{\ell_2}=1}\|Su\|_{\ell_2}^2 \in [0,1].$$

If the columns of S are orthogonal, then $\nu(S) = 1$; if the columns of S are strongly aligned, then $\nu(S) \approx 0$. Given noisy data, finding all of the directions in the range of S is harder when the permeance is small.

7.3. **Gross errors in matrix entries.** In this section, we consider a noise model in which a moderate number of entries of the data matrix are corrupted arbitrarily. Suppose that we observe

$$B = L_0 + \Omega_0 \in \mathbb{R}^{n\times m}, \tag{7.1}$$

where L_0 is a low-rank data matrix, while Ω_0 models contamination of individual entries. We imagine that Ω_0 is sparse, but its nonzero entries are arbitrary. In particular, the noise can erase individual entries of the data matrix or change them maliciously.

The data model (7.1) was first studied by Chandrasekaran et al. [CSPW11]. Given the observation B, they proposed to separate the low-rank data matrix L_0 from the noise Ω_0 by solving a semidefinite optimization problem:

$$\underset{L,\Omega\in\mathbb{R}^{n\times m}}{\text{minimize}} \quad \|L\|_{S_1} + \frac{1}{\sqrt{\max\{n,m\}}}\|\Omega\|_{\ell_1} \quad \text{subject to} \quad B = L + \Omega. \tag{7.2}$$

The Schatten 1-norm $\|\cdot\|_{S_1}$ promotes low-rank in the component L, while the vectorized ℓ_1 norm $\|\cdot\|_{\ell_1}$ promotes sparsity in the other component Ω.

Chandrasekaran et al. [CSPW11] developed deterministic conditions under which the pair (L_0, Ω_0) is the unique solution of the optimization problem (7.2). The subsequent paper [CLMW11] established additional guarantees under a probabilistic model for the data. By adapting these arguments, we can show that it is often possible to remove sparse errors from a realization of the Gaussian loadings model. We have the following result.

Theorem 7.2 (Gaussian loadings model: Sparse noise)**.** *Suppose we observe the matrix $B = L_0 + \Omega_0 \in \mathbb{R}^{n\times m}$ where*

    (1) The sign matrix $S \in \{\pm 1\}^{n\times r}$ is Schur independent, with permeance $\nu(S)$;
    (2) The data matrix $L_0 \sim \mathrm{GLM}(S, m)$ is drawn from the Gaussian loadings model; and
    (3) The noise matrix Ω_0 has ω nonzero entries, in uniformly random locations, with arbitrary magnitudes that may depend on L_0.

Assume that the parameters satisfy

$$\nu(S) \gtrsim \frac{r\log^3(n+m)}{\min\{n,m\}} \quad \text{and} \quad \omega \lesssim mn. \tag{7.3}$$

Then, with high probability over the randomness in the model, the pair (L_0, Ω_0) is the unique solution to the optimization problem (7.2).

See Appendix A for the proof, which relies on standard methods from high-dimensional probability. Better results may be possible with additional argument.

Let us take a moment to discuss Theorem 7.2. The condition (7.3) on the sparsity ω of the noise allows us to repair errors in a constant proportion of the entries of the data matrix. The condition (7.3) on the permeance $\nu(S)$ states that the conditioning of the sign matrix S controls the minimum number m of samples that we need to remove the noise. It also gives a minimum standard on the conditioning, in terms of the dimension n, for the theorem to operate.

7.4. **Gross errors in matrix columns.** In this section, we study a noise model in which a moderate number of columns of the data matrix are corrupted arbitrarily. Suppose that we observe a matrix of the form

$$B = \begin{bmatrix} L_0 & \Omega_0 \end{bmatrix} \Pi \in \mathbb{R}^{n \times M} \tag{7.4}$$

where $L_0 \in \mathbb{R}^{n \times m}$ is a rank-r matrix of clean data, $\Omega_0 \in \mathbb{R}^{n \times m'}$ contains arbitrary noise, $\Pi \in \mathbb{R}^{M \times M}$ is an unknown permutation, and $M = m + m'$. We regard the columns of Ω_0 as outliers that are mixed in with the data and that we need to remove. The number of inliers (m) and outliers (m') is not known in advance, but we do require the value r of the rank.

The data model (7.4) has a long history in robust statistics [HR09]. In recent years, researchers have proposed a number of methods for removing outliers by means of convex optimization. We outline a technique that is based on the following intuition [LMTZ15]. Observe that the orthogonal projector $P_0 \in \mathbb{H}_n$ onto the range of L_0 discriminates inliers from outliers. Indeed, $P_0 x = x$ for $x \in \mathrm{range}(L_0)$, while $P_0 x \neq x$ for $x \notin \mathrm{range}(L_0)$. Therefore, we may try to find a rank-r projector that fixes many columns of B by charging as much as possible for columns that are not reproduced.

If we can determine the projector P_0, then we can find the low-rank matrix L_0 by picking out the columns of B that are fixed by the projector P_0. (It is possible that P_0 also fixes some columns of the noise matrix Ω_0, but then we should probably regard these outliers as inliers.)

Lerman et al. [LMTZ15] propose to find the projector P_0 by solving the semidefinite optimization problem

$$\underset{P \in \mathbb{H}_n}{\text{minimize}} \quad \sum_{i=1}^{M} \|(I - P)B e_i\|_{\ell_2} \quad \text{subject to} \quad \mathrm{trace}(P) = r \text{ and } 0 \preccurlyeq P \preccurlyeq I. \tag{7.5}$$

The constraint set is the best convex relaxation of the set of rank-r orthogonal projectors.

Lerman et al. [LMTZ15] develop deterministic conditions under which the orthogonal projector P_0 onto $\mathrm{range}(L_0)$ is the unique solution to the problem (7.5), and they specialize these results to random data models. By adapting their arguments, we can establish that it is possible to remove outliers from a data matrix that follows the Gaussian loadings model.

Theorem 7.3 (Gaussian loadings model: Outliers). *Suppose we observe the matrix* $B = \begin{bmatrix} L_0 & \Omega_0 \end{bmatrix} \Pi \in \mathbb{R}^{n \times M}$, *where* $M = m + m'$ *and*

- *(1) The sign matrix $S \in \{\pm 1\}^{n \times r}$ is Schur independent, with permeance $\nu(S)$;*
- *(2) The data matrix $L_0 \sim \mathrm{GLM}(S, m)$ is drawn from the Gaussian loadings model;*
- *(3) The noise matrix $\Omega_0 \in \mathbb{R}^{n \times m'}$ is standard normal; and*
- *(4) The permutation matrix $\Pi \in \mathbb{R}^{M \times M}$ is arbitrary.*

Assume that the parameters satisfy

$$\max\left\{1, \frac{m'}{n}\right\} \lesssim \sqrt{\nu(S)} \cdot \frac{m}{r}. \tag{7.6}$$

Then, with high probability over the randomness in the model, the orthogonal projector onto $\mathrm{range}(L_0)$ *is the unique solution to the optimization problem* (7.2).

The proof of Theorem 7.3 appears in Appendix B. Results for other data models are also possible.

Let us discuss Theorem 7.3 briefly. First, under the data model, the columns of L_0 and Ω_0 all have the same energy, on average, so we cannot distinguish the inliers and the outliers on the basis of their norm. Next, the condition in (7.6) requires that the number m' of outliers relative to the ambient dimension n is no greater than the number m of inliers relative to the dimension r of the space spanned by the inliers. The permeance $\nu(S)$ also affects how many outliers we can tolerate; it is easier to reject outliers when the columns of L_0 are well distributed.

7.5. **Future work.** Our results on robustness (Theorem 7.2 and Theorem 7.3) are based on the fact that subspaces have a very strong signal, so it is easy to reject noise that violates the subspace structure. We can exploit the same intuition to handle some other types of noise models. On the other hand, these

ideas do not suffice to treat robustness of the sign component decomposition with respect to, say, additive Gaussian noise. An important direction for future research is to develop additional tools for producing a sign component decomposition of a noisy data matrix drawn from a more general model.

APPENDIX A. GAUSSIAN LOADINGS MODEL: REMOVING GROSS ERRORS IN MATRIX ENTRIES

A.1. **Background.** We consider a noise model in which a moderate number of entries of a data matrix are corrupted arbitrarily. We observe

$$\boldsymbol{B} = \boldsymbol{L}_0 + \boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m},$$

where $\boldsymbol{L}_0 \in \mathbb{R}^{n \times m}$ is a low-rank data matrix and $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m}$ models contamination of individual entries. We assume that $\boldsymbol{\Omega}_0$ is sparse and both contributions to $\boldsymbol{B}$ are well separated in the following sense. The low rank data matrix $\boldsymbol{L}_0$ cannot be too sparse, while the sparse noise corruption $\boldsymbol{\Omega}_0$ cannot be too low-rank. Choosing the support of $\boldsymbol{\Omega}_0$ uniformly at random takes care of the latter condition. The *incoherence statistics* has been identified as a witness for the former. Set $r = \mathrm{rank}(\boldsymbol{L}_0)$ and let $\boldsymbol{L}_0 = \boldsymbol{U}\boldsymbol{\Sigma}\boldsymbol{V}^{\mathrm{t}}$ be a SVD (2.2). Note that the orthonormal matrices $\boldsymbol{U} \in \mathbb{R}^{n \times r}$ and $\boldsymbol{V} \in \mathbb{R}^{m \times r}$ are closely related to the projectors onto row- and column range of $\boldsymbol{L}_0$. More precisely, $\boldsymbol{P}_{\mathrm{ran}(\boldsymbol{L}_0)} = \boldsymbol{U}\boldsymbol{U}^{\mathrm{t}} \in \mathbb{H}_n$ and $\boldsymbol{P}_{\mathrm{ran}(\boldsymbol{L}_0^{\mathrm{t}})} = \boldsymbol{V}\boldsymbol{V}^{\mathrm{t}} \in \mathbb{H}_m$. Based on the SVD, we define three *incoherence parameters*:

$$\begin{aligned}
\mu_{(l)}(\boldsymbol{L}_0) &= \frac{n}{r} \max_{1 \le k \le n} \left\| \boldsymbol{P}_{\mathrm{ran}(\boldsymbol{L}_0)} \mathbf{e}_k \right\|_{\ell_2}^2 = \frac{n}{r} \max_{1 \le k \le n} \left\| \boldsymbol{U}^{\mathrm{t}} \mathbf{e}_k \right\|_{\ell_2}^2, \\
\mu_{(r)}(\boldsymbol{L}_0) &= \frac{m}{r} \max_{1 \le l \le m} \left\| \boldsymbol{P}_{\mathrm{ran}(\boldsymbol{L}_0^{\mathrm{t}})} \mathbf{e}_l \right\|_{\ell_2}^2 = \frac{m}{r} \max_{1 \le l \le m} \left\| \boldsymbol{V}^{\mathrm{t}} \mathbf{e}_l \right\|_{\ell_2}^2, \\
\tilde{\mu}(\boldsymbol{L}_0) &= \frac{nm}{r} \max_{1 \le k \le n, 1 \le l \le m} \left| \langle \mathbf{e}_k, \boldsymbol{U}\boldsymbol{V}^{\mathrm{t}} \mathbf{e}_l \rangle \right|^2.
\end{aligned} \tag{A.1}$$

The first two parameters have a ready explanation: they measure how well spread-out the left- and right-singular vectors are with respect to the standard basis. Small values ensure that these vectors are not (too) sparse. The third parameter lacks a compelling interpretation. It should be viewed as a technical requirement that features prominently in the arguments by Candès et al. [CLMW11]. Define the maximum of these three parameters

$$\mu(\boldsymbol{L}_0) = \max\left\{ \mu_{(l)}(\boldsymbol{L}_0), \mu_{(r)}(\boldsymbol{L}_0), \tilde{\mu}(\boldsymbol{L}_0) \right\} \tag{A.2}$$

and consider the following semidefinite program for rank-sparsity decomposition [CSPW11]:

$$\operatorname*{minimize}_{\boldsymbol{L}, \boldsymbol{\Omega} \in \mathbb{R}^{n \times m}} \quad \|\boldsymbol{L}\|_{S_1} + \frac{1}{\sqrt{\max\{n,m\}}} \|\boldsymbol{\Omega}\|_{\ell_1} \quad \text{subject to} \quad \boldsymbol{B} = \boldsymbol{L} + \boldsymbol{\Omega}. \tag{A.3}$$

The main result [CLMW11, Thm. 1] establishes a probabilistic recovery guarantee for rank-sparsity decomposition.

Theorem A.1 (Candès, Li, Ma & Wright). *Suppose that $\boldsymbol{B} = \boldsymbol{L}_0 + \boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m}$ is such that*

(1) $\boldsymbol{L}_0$ is such that rank $r = \mathrm{rank}(\boldsymbol{L}_0)$ and incoherence $\mu(\boldsymbol{L}_0)$ (A.2) obey

$$r \le \rho_{(r)} \frac{\min\{n, m\}}{\mu(\boldsymbol{L}_0) \log^2(n + m)};$$

(2) $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m}$ has ω nonzero entries, in uniformly random locations, with arbitrary magnitudes that may depend on $\boldsymbol{L}_0$.

Then, with probability at least $1 - \gamma \max\{n, m\}^{-10}$ (over the choice of the sparsity pattern), the solution to the semidefinite program (A.3) is $(\boldsymbol{L}_\star, \boldsymbol{\Omega}_\star) = (\boldsymbol{L}_0, \boldsymbol{\Omega}_0)$. Here, $\rho_{(r)}, \rho_{(s)}$ and γ are constants of appropriate size.

A.2. **Coherence statistics for the Gaussian loadings model and denoising.** Recall the *Gaussian loadings model* (Model 7.1). Select a Schur-independent sign matrix $S \in \{\pm 1\}^{n \times r}$ and a number $m \geq r$. A matrix $L_0 \sim \text{GLM}(S, m)$ takes the form $L_0 = r^{-1/2} S G^{\mathrm{t}}$, where $G \in \mathbb{R}^{m \times r}$ is standard normal. The *permeance statistic* captures how well the columns of L_0 fill out their span:

$$\nu(S) = \frac{1}{n} \lambda_{\min}(S^{\mathrm{t}} S) = \frac{1}{n} \min_{\|u\|_{\ell_2} = 1} \|S u\|_{\ell_2}^2. \tag{A.4}$$

The following technical statement implies that the permeance also controls two of the three incoherence parameters (A.1) associated with the Gaussian loadings model.

Lemma A.2. *Fix $S \in \{\pm 1\}^{n \times r}$ with permeance $\nu(S) > 0$ and a natural number $m \geq r$. Choose L_0 from the Gaussian loadings model. Then, for any $\alpha > 0$,*

$$\mu_{(l)}(L_0) \leq \nu(S)^{-1},$$
$$\Pr\left[\mu_{(r)}(L_0) \geq 4(\alpha + 1)\log(n + m)\right] \leq 2(nm)^{-\alpha}/n,$$
$$\Pr\left[\tilde{\mu}(L_0) \geq \tfrac{4(\alpha+1)}{\nu(S)}\log(n + m)\right] \leq 2(nm)^{-\alpha}.$$

The first bound is deterministic and original. The remaining probabilistic bounds follow from adapting arguments by Candès et al. [CLMW11, Sec. 2.2] to the Gaussian loadings model. Inserting these bounds into [CLMW11, Thm. 1] readily implies that gross errors in entries of a Gaussian loadings matrix can be removed completely.

Theorem A.3 (Theorem 7.2, restatement)**.** *Suppose we observe the matrix $B = L_0 + \Omega_0 \in \mathbb{R}^{n \times m}$ where*

(1) *The sign matrix $S \in \{\pm 1\}^{n \times r}$ is Schur independent, with permeance $\nu(S)$;*
(2) *The data matrix $L_0 \sim \text{GLM}(S, m)$ is drawn from the Gaussian loadings model; and*
(3) *The noise matrix Ω_0 has ω nonzero entries, in uniformly random locations, with arbitrary magnitudes that may depend on L_0.*

Assume that the parameters satisfy $\omega \leq mn$ and

$$\nu(S) \gtrsim \frac{r \log^3(n + m)}{\min\{n, m\}} \tag{A.5}$$

Then, with high probability over the randomness in the model, the pair (L_0, Ω_0) is the unique solution to the optimization problem (A.3).

Proof. Sample L_0 from the Gaussian loadings model $\text{GLM}(S, m)$. Choose $\alpha = 1$ and combine Lemma A.2 with a union bound to conclude

$$\mu(L_0) = \max\left\{\mu_{(r)}(L_0), \mu_{(l)}(L_0), \tilde{\mu}(L_0)\right\} \leq \tfrac{8}{\nu(S)}\log(n + m)$$

with probability at least $1 - \frac{4}{nm}$. Conditioned on this event, condition (1) in Theorem A.1 becomes equivalent to (A.5) up to constants. Moreover, condition (2) is met by assumption. Apply Theorem A.1 to complete the argument.   □

A.3. **Proof of Lemma A.2.** We now proceed to proving Lemma A.2 and address the three different bounds separately. Suppose that $L_0 = S G^{\mathrm{t}} \sim \text{GLM}(S, m)$ is sampled from the Gaussian loadings model. Then, $\nu(S)$ ensures that S has full column rank, while $m \geq r$ implies that the standard normal matrix G has full column rank with probability one. This ensures that both left- and right incoherence are determined by the individual factors:

$$\mu_{(l)}(L_0) = \mu_{(l)}(S) \quad \text{and} \quad \mu_{(r)}(L_0) = \mu_{(r)}(G). \tag{A.6}$$

A.3.1. *Deterministic bound for the left-incoherence.*

Lemma A.4. *A sign matrix $S \in \{\pm 1\}^{n \times r}$ with permeance $\nu(S) > 0$ obeys*

$$\mu_{(l)}(S) \leq \nu(S)^{-1}.$$

Proof. The function $f : \mathbb{H}_n \to \mathbb{R}$ defined by $X \mapsto \max_{1 \leq k \leq n} \langle \mathbf{e}_k, X\mathbf{e}_k \rangle$ is monotone with respect to the psd order, i.e. $X \succcurlyeq Y$ implies $f(X) \geq f(Y)$. Next, let $P \in \mathbb{H}_n$ denote the projector onto the range of S. Then, $P \preccurlyeq \lambda_{\min \neq 0}^{-1} SS^{\mathsf{t}}$, where $\lambda_{\min \neq 0}(SS^{\mathsf{t}})$ is the smallest non-zero eigenvalue of SS^{t}. Apply an SVD $S = U\Sigma V^{\mathsf{t}}$ and use orthogonal invariance of eigenvalues to conclude

$$\lambda_{\min \neq 0}(SS^{\mathsf{t}}) = \lambda_{\min \neq 0}\left(U\Sigma^2 U^{\mathsf{t}}\right) = \sigma_r^2 = \lambda_{\min}\left(V\Sigma^2 V^{\mathsf{t}}\right) = \lambda_{\min}\left(S^{\mathsf{t}}S\right) = n\nu(S).$$

Combine this observation with monotonicity of $f : \mathbb{H}_n \to \mathbb{R}$:

$$\mu_{(l)}(S) = \frac{n}{r} \max_{1 \leq k \leq n} \|P\mathbf{e}_k\|_{\ell_2}^2 = \frac{n}{r} f(P) \leq \frac{n}{r} f\left(\tfrac{1}{n\nu(S)} SS^{\mathsf{t}}\right) = \frac{1}{r\nu(S)} \max_{1 \leq k \leq n} \langle \mathbf{e}_k, SS^{\mathsf{t}}\mathbf{e}_k \rangle.$$

Finally, note that any sign matrix $S = \begin{bmatrix} s_1 & \dots & s_r \end{bmatrix} \in \{\pm 1\}^{n \times r}$ obeys $\langle \mathbf{e}_k, SS^{\mathsf{t}}\mathbf{e}_k \rangle = \sum_{i=1}^r \langle \mathbf{e}_k, s_i \rangle^2 = r$ for all $1 \leq k \leq n$ simultaneously. □

A.3.2. *Probabilistic bound for the right incoherence.* Rel. (A.6) asserts that the right-incoherence of a Gaussian loadings sample is fully characterized by the standard normal matrix $G \in \mathbb{R}^{m \times r}$. Rotation invariance of the standard normal columns extends to the range of G. Moreover, with probability one, this range is r-dimensional. Condition on this almost sure event. Then, the projector $Q \in \mathbb{H}_m$ on the range of G is a random matrix of the form

$$Q = R\left(\sum_{i=1}^r \mathbf{e}_i\mathbf{e}_i^{\mathsf{t}}\right) R^{\mathsf{t}} \quad \text{where} \quad R \stackrel{\text{unif.}}{\sim} \mathsf{O}(m). \tag{A.7}$$

Here $R \stackrel{\text{unif.}}{\sim} \mathsf{O}(m)$ implies that the matrix R is chosen with respect to the unique invariant measure on the orthogonal group. This measure is also known as the Haar measure. For technical reasons, we point out a reformulation of Eq. (A.7). The orthogonal group is a compact Lie group and therefore unimodular. This implies that the Haar measure is invariant under taking inverses (transpositions):

$$Q = R^{\mathsf{t}}\left(\sum_{i=1}^r \mathbf{e}_i\mathbf{e}_i^{\mathsf{t}}\right) R \quad \text{where} \quad R \stackrel{\text{unif.}}{\sim} \mathsf{O}(m). \tag{A.8}$$

Lemma A.5. *Let $G \in \mathbb{R}^{m \times r}$ be a random rank-r matrix with rotation-invariant range in the sense of Eq. (A.8). Then, for any $\alpha > 0$*

$$\Pr\left[\mu_{(l)}(G) \geq 4(\alpha + 1)\log(n + m)\right] \leq 2(nm)^{-\alpha}/n.$$

This result covers Gaussian random matrices, as well as the random orthogonal model discussed in Ref. [CLMW11].

Proof. Fix a standard basis vector $\mathbf{e}_{l_0} \in \mathbb{R}^m$ ($1 \leq l_0 \leq m$) and consider the random variable $\langle \mathbf{e}_{l_0}, Q\mathbf{e}_{l_0} \rangle$. Use rotation invariance (A.8) to reformulate the distribution of this random variable:

$$\langle \mathbf{e}_{l_0}, Q\mathbf{e}_{l_0} \rangle = \langle R\mathbf{e}_{l_0}, \left(\sum_{i=1}^r \mathbf{e}_i\mathbf{e}_i^{\mathsf{t}}\right) R\mathbf{e}_{l_0} \rangle = \sum_{i=1}^r \langle \mathbf{e}_i, R\mathbf{e}_{l_0} \rangle^2 \quad \text{where} \quad R \stackrel{\text{unif}}{\sim} \mathsf{O}(m).$$

The uniform distribution of R implies that the unit vector $\hat{v} = R\mathbf{e}_{l_0}$ is distributed uniformly over the unit sphere in $\mathbb{R}^m$. Each component $\langle \mathbf{e}_i, v \rangle$ is approximately normal with mean zero and variance $m^{-1/2}$. This suggests that $\langle \mathbf{e}_{l_0}, Q\mathbf{e}_{l_0} \rangle$ resembles a re-normalized χ^2-distribution with r degrees of freedom. Lemma A.7 below makes this intuition precise and ensures

$$\Pr\left[\sum_{i=1}^r \langle e_i, \hat{v} \rangle^2 \geq \tfrac{4(\alpha+1)r}{m}\log(n + m)\right] \leq 2(nm)^{-(\alpha+1)}.$$

The final claim follows from a union bound over all m possible choices of the standard basis vector e_{l_0}. □

A.3.3. *Probabilistic bound for the third incoherence parameter.*

Lemma A.6. *Fix $S \in \{\pm 1\}^{n \times r}$ with permeance $\nu(S) > 0$, $m \geq r$ and sample $L_0 \sim \mathrm{GLM}(S, m)$ from the Gaussian loadings model. Then, for any $\alpha > 0$*

$$\Pr\left[\tilde{\mu}(L_0) \geq \tfrac{4(\alpha+1)}{\nu(S)} \log(n+m)\right] \leq 2(nm)^{-\alpha}.$$

Proof. Let $L_0 = U\Sigma V^{\mathrm{t}}$ be a SVD. By definition of $\tilde{\mu}(\cdot)$, we must establish

$$\Pr\left[\max_{1 \leq k \leq n, 1 \leq l \leq m} \left|\langle \mathbf{e}_l, VU^{\mathrm{t}}\mathbf{e}_k\rangle\right|^2 \geq \tfrac{4(\alpha+1)r}{\nu(S)mn} \log(n+m)\right] \leq 2(nm)^{-\alpha}.$$

Fix $e_{k_0} \in \mathbb{R}^n$ and $e_{l_0} \in \mathbb{R}^m$. Set $u = U^{\mathrm{t}} e_{k_0} \in \mathbb{R}^r$ and note that its Euclidean length is controlled via left-incoherence and Lemma A.4:

$$\|U^{\mathrm{t}} e_{k_0}\|_{\ell_2}^2 \leq \max_{1 \leq k \leq n} \|U^{\mathrm{t}} e_k\|_{\ell_2}^2 = \frac{r}{n}\mu_{(l)}(L_0) = \frac{r}{n}\mu_{(l)}(S) \leq \frac{r}{n\nu(S)}.$$

Let $\hat{u} = U^{\mathrm{t}} e_{k_0}/\|U^{\mathrm{t}} e_{k_0}\|_{\ell_2}$ be the unit vector pointing into the same direction as u. Then, this bound on the Euclidean length ensures

$$\Pr\left[\left|\langle e_{l_0}, VU^{\mathrm{t}} e_{k_0}\rangle\right|^2 \geq \tfrac{4(\alpha+1)r}{\nu(S)nm} \log(n+m)\right] = \Pr\left[\left|\langle e_{l_0}, V\hat{u}\rangle\right|^2 \geq \tfrac{4(\alpha+1)r}{\nu(S)nm} \|U^{\mathrm{t}} e_{k_0}\|_{\ell_2}^{-2} \log(n+m)\right]$$

$$\leq \Pr\left[\left|\langle e_{l_0}, V\hat{u}\rangle\right|^2 \geq \tfrac{4(\alpha+1)}{m} \log(n+m)\right]$$

The vector $V\hat{u} \in \mathbb{R}^m$ also has unit length, because V is orthonormal. What is more, rotational invariance of the Gaussian matrix G extends to the matrix of right singular vectors: $V \colon V \sim RV$ for any $R \in \mathrm{O}(m)$. This implies that $V\hat{u}$ is distributed uniformly on the m-dimensional unit sphere. Lemma A.7 then implies

$$\Pr\left[|\langle e_{l_0}, V\hat{u}\rangle|^2 \geq \tfrac{4(\alpha+1)}{m} \log(n+m)\right] \leq 2(nm)^{-(\alpha+1)}.$$

Finally, deduce the claim by applying a union bound over all nm choices of standard basis vectors $\mathbf{e}_{k_0} \in \mathbb{R}^n$ and $\mathbf{e}_{l_0} \in \mathbb{R}^m$. □

A.4. **Exponential tail bound for vectors chosen uniformly from the unit sphere.**

Lemma A.7. *Let $\{u_1, \ldots, u_r\} \subset \mathbb{R}^m$ be a set of unit vectors and choose $\hat{v}$ uniformly from the unit sphere in m dimensions. Then, for any $t > 0$*

$$\Pr\left[\sum_{i=1}^r \langle u_i, \hat{v}\rangle^2 \geq \tfrac{4rt}{m}\right] \leq 2\mathrm{e}^{-t}.$$

This is a typical tail bound. It follows from establishing sub-exponential moment growth and subsequently applying the exponential Markov inequality. Stronger concentration inequalities follow from concentration of measure. However, these tighter bounds are somewhat unwieldy by comparison.

Proof. Define the random variable $S = \frac{m}{r}\sum_{i=1}^r \langle u_i, \hat{v}\rangle^2$ and interpret it as the squared Euclidean norm of a random vector $\tilde{v} = \sqrt{\frac{m}{r}}\left[\langle u_1, \hat{v}\rangle \quad \ldots \quad \langle u_1, \hat{v}\rangle\right]^{\mathrm{t}} \in \mathbb{R}^r$. The fundamental ℓ_p-norm relation relation $\|\tilde{v}\|_{\ell_2} \leq r^{1/2-1/(2p)}\|\tilde{v}\|_{\ell_{2p}}$ in $\mathbb{R}^r$ then implies

$$\mathbb{E}\left[S^p\right] = \mathbb{E}\left[\|\tilde{v}\|_{\ell_2}^{2p}\right] \leq r^{p-1}\mathbb{E}\left[\|\tilde{v}\|_{\ell_{2p}}^{2p}\right] = \frac{m^p}{r}\sum_{i=1}^r \mathbb{E}\left[\langle u_i, v\rangle^{2p}\right] \quad \text{for any} \quad p \in \mathbb{N}.$$

Next, let $g \in \mathbb{R}^m$ be a standard normal vector. This random vector may be decomposed into a direction $\hat{v} = g/\|v\|_{\ell_2}$ and a radius $\rho = \|g\|_{\ell_2}$. The direction is distributed uniformly on the m-dimensional unit sphere and ρ^2 follows a χ^2-distribution with m degrees of freedom. Importantly, $\hat{v}$ and ρ are stochastically independent. Combine this with rotation invariance and normalization ($\|u_i\|_{\ell_2} = 1$) to conclude

$$\mathbb{E}\left[\langle u_i, \hat{v}\rangle^{2p}\right] = \frac{\mathbb{E}\left[\langle u_i, g\rangle^{2p}\right]}{\mathbb{E}\left[\|g\|_{\ell_2}^{2p}\right]} = \frac{\|u_i\|_{\ell_2}^{2p}\mathbb{E}\left[\langle \mathbf{e}_1, g\rangle^{2p}\right]}{\mathbb{E}\left[(\chi_m^2)^p\right]} = \frac{(2p-1)!!}{m(m+2)\cdots(m+2(p-1))} \quad \text{for all} \quad 1 \leq i \leq r.$$

The last equation follows from well-known expressions for the moments of standard normal and χ_m^2-distributed random variables. The relations $(2p-1)!! \le (2p)!! = 2^p p!$ and $m(m+2)\cdots(m+2(p-1)) \ge m^p$ then imply

$$\mathbb{E}\left[S^p\right] \le \frac{m^p}{r}\sum_{i=1}^r \frac{2^p p!}{m^p} = 2^p p! \quad \text{for all} \quad p \in \mathbb{N}.$$

This moment growth suggests sub-exponential tail behavior. The exponential Markov inequality makes this intuition precise:

$$\Pr\left[\sum_{i=1}^r \langle \boldsymbol{u}_i, \hat{\boldsymbol{v}}\rangle^2 \ge \tfrac{4rt}{m}\right] = \Pr\left[\tfrac{S}{4} \ge t\right] = \Pr\left[e^{S/4} \ge e^t\right] \le e^{-t}\mathbb{E}\left[\exp\left(\tfrac{S}{4}\right)\right]$$

$$= e^{-t}\sum_{p=0}^\infty \frac{\mathbb{E}[S^p]}{4^p p!} \le e^{-t}\sum_{p=0}^\infty 2^{-p} = 2e^{-t}.$$

$\square$

APPENDIX B. GAUSSIAN LOADINGS MODEL: REMOVING GROSS ERRORS IN MATRIX COLUMNS

B.1. **Background.** We consider a noise model in which a data matrix is corrupted by a potentially large number of outliers [HR09, LMTZ15]. Consider a data matrix that consists of m low-dimensional inliers and m' high-dimensional outliers. More precisely, set $M = m + m'$ and consider a matrix of the form

$$\boldsymbol{B} = \begin{bmatrix} \boldsymbol{L}_0 & \boldsymbol{\Omega}_0 \end{bmatrix} \boldsymbol{\Pi} \in \mathbb{R}^{n \times M}.$$

Here, $\boldsymbol{L}_0 \in \mathbb{R}^{n \times m}$ is a rank-r matrix of clean data, $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m'}$ subsumes arbitrary noise and $\boldsymbol{\Pi} \in \mathbb{R}^{M \times M}$ is an unknown permutation. We do not know m (number of inliers) and m' (number of column outliers) in advance, but we require knowledge of $r = \text{rank}(\boldsymbol{L}_0)$. This In&Out model has a key feature. The inliers (columns of $\boldsymbol{L}_0$) are confined to the r-dimensional range $\mathsf{L}_0 \subset \mathbb{R}^n$ of $\boldsymbol{L}_0$, while the outliers can be arbitrary. Hence, the projector $\boldsymbol{P}_0 \in \mathbb{H}_n$ onto L_0 discriminates inliers from outliers. The *reaper* algorithm [LMTZ15] is designed to recover this projector from the corrupted data matrix:

$$\underset{\boldsymbol{P} \in \mathbb{H}_n}{\text{minimize}} \quad \sum_{i=1}^M \|(\mathbf{I} - \boldsymbol{P})\boldsymbol{B}\mathbf{e}_i\|_{\ell_2} \quad \text{subject to} \quad \text{trace}(\boldsymbol{P}) = r \text{ and } \mathbf{0} \preccurlyeq \boldsymbol{P} \preccurlyeq \mathbf{I}. \tag{B.1}$$

Sufficient conditions for exact recovery have been established by Lerman, McCoy, Tropp and Zhang [LMTZ15]. These criteria are deterministic and depend on data-dependent parameters that play a role analogous to the incoherence statistics discussed in Appendix A. Define the *permeance statistics* for the clean data matrix $\boldsymbol{L}_0 = \begin{bmatrix} \boldsymbol{l}_1 & \dots & \boldsymbol{l}_m \end{bmatrix}$ with range L_0:

$$P(\mathsf{L}_0, \boldsymbol{L}_0) = \inf_{\boldsymbol{u} \in \mathsf{L}_0, \|\boldsymbol{u}\|_{\ell_2}=1} \sum_{i=1}^m |\langle \boldsymbol{u}, \boldsymbol{l}_i\rangle|. \tag{B.2}$$

Remark B.1 (Relation between permeance statistics and permeance parameter). Permeance parameter (A.4) and permeance statistics (B.2) are closely related. The permeance statistics minimizes a ℓ_1-penalty, while the permeance parameter optimizes the least-squares loss of the same objective:

$$n\nu(\boldsymbol{S}) = \lambda_{\min}(\boldsymbol{S}^\mathsf{t}\boldsymbol{S}) = \inf_{\boldsymbol{u} \in \text{ran}(\boldsymbol{S}), \|\boldsymbol{u}\|_{\ell_2}=1} \|\boldsymbol{S}\boldsymbol{u}\|_{\ell_2}^2 = \inf_{\boldsymbol{u} \in \text{ran}(\boldsymbol{S}), \|\boldsymbol{u}\|_{\ell_2}=1} \sum_{i=1}^r \langle \boldsymbol{u}, \boldsymbol{s}_i\rangle^2.$$

Define the *spherical linear structure statistics* for the matrix of outliers $\boldsymbol{\Omega}_0 = \begin{bmatrix} \boldsymbol{\omega}_1 & \dots & \boldsymbol{\omega}_{m'} \end{bmatrix}$ and the ortho-complement $\mathsf{L}_0^\perp$ of L_0:

$$\hat{S}\left(\mathsf{L}_0^\perp, \boldsymbol{\Omega}_0\right)^2 = \sup_{\boldsymbol{u} \in \mathsf{L}_0^\perp, \|\boldsymbol{u}\|_{\ell_2}=1} \sum_{i=1}^{m'} \left|\langle \boldsymbol{u}, \widehat{\boldsymbol{P}_{\mathsf{L}_0^\perp}\boldsymbol{\omega}_i}\rangle\right|^2.$$

Here, $\boldsymbol{P}_{\mathsf{L}_0^\perp} \in \mathbb{H}_n$ denotes the orthogonal projector onto $\mathsf{L}_0^\perp$ and $\hat{\boldsymbol{x}} = \boldsymbol{x}/\|\boldsymbol{x}\|_{\ell_2}$ denotes the unit vector pointing into the same direction as $\boldsymbol{x}$ (with the convention that $\hat{\mathbf{0}} = \mathbf{0}$). Finally, let $\|\cdot\|_{S_\infty}$ denote the spectral (or operator) norm. The main result [LMTZ15, Thm. 3.1] establishes a deterministic recovery guarantee for the reaper algorithm (B.1) based on these parameters.

Theorem B.2 (Lerman, McCoy, Tropp & Zhang). *Consider a data matrix $\boldsymbol{B} = \begin{bmatrix} \boldsymbol{L}_0 & \boldsymbol{\Omega}_0 \end{bmatrix} \boldsymbol{\Pi} \in \mathbb{R}^{n \times M}$ and set* $\mathsf{L}_0 = \mathrm{ran}(\boldsymbol{L}_0) \subset \mathbb{R}^n$. *Suppose that*

$$P(\mathsf{L}_0, \boldsymbol{L}_0) > \sqrt{2r}\,\hat{S}\big(\mathsf{L}_0^\perp, \boldsymbol{\Omega}_0\big) \|\boldsymbol{\Omega}_0\|_{S_\infty}.$$

Then, the solution to the reaper problem (B.1) *is the orthogonal projection $\boldsymbol{P}_0$ onto* L_0.

B.2. **Summary statistics parameters for the Gaussian loadings model and denoising.** We consider a clean data matrix sampled from the Gaussian loadings model (Model 7.1):

$$\boldsymbol{L}_0 = \frac{1}{\sqrt{r}} \boldsymbol{S} \boldsymbol{G}^{\mathsf{t}} \quad \text{where} \quad \boldsymbol{S} \in \{\pm 1\}^{n \times r} \quad \text{and} \quad \boldsymbol{G} \in \mathbb{R}^{m \times r} \quad \text{has standard normal entries.}$$

The relation between permeance statistics and permeance parameter displayed in Remark B.1 suggests that the former is controlled by the square root of the latter. The following technical results makes this intuition precise.

Lemma B.3. *Let $\mathsf{L}_0 \in \mathbb{R}^n$ denote the range of a sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ with permeance $\nu(\boldsymbol{S})$. A sample $\boldsymbol{L}_0 \sim \mathrm{GLM}(\boldsymbol{S}, m)$ from the associated Gaussian loadings model obeys*

$$\Pr\left[P(\mathsf{L}_0, \boldsymbol{L}_0) \geq \frac{1}{6}\sqrt{\frac{nm}{r}}\left(\sqrt{m\nu(\boldsymbol{S})} - 12\sqrt{r} - 2\sqrt{\nu(\boldsymbol{S})}\,t\right)\right] \geq 1 - \mathrm{e}^{-t^2/2} \quad \text{for any} \quad t > 0.$$

This claim follows from a versatile proof technique known as Mendelson's small ball method [Men15, KM15, Tro15]. We present a detailed proof in the next sub-section.

Next, we consider the impact of white noise corruptions. Let $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m'}$ be a random matrix with standard normal entries. The spectral properties of such random matrices are very well understood:

$$\Pr\left[\|\boldsymbol{\Omega}_0\|_{S_\infty} \leq \sqrt{n} + \sqrt{m'} + t \right] \geq 1 - \mathrm{e}^{-t^2} \quad \text{for any} \quad t > 0, \tag{B.3}$$

see e.g. [DS01, Theorem 2.13]. Similar results remain valid for random matrices whose columns are sampled uniformly from the m'-dimensional unit sphere instead. We borrow the following extension from Lerman et al.

Lemma B.4 (Lem. 8.4 in [LMTZ15]). *Fix a r-dimensional subspace $\mathsf{L}_0 \in \mathbb{R}^n$. Then, a standard Gaussian matrix $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m'}$ obeys*

$$\Pr\left[\hat{S}\big(\mathsf{L}_0^\perp, \boldsymbol{\Omega}_0\big) \leq \frac{\sqrt{m'} + \sqrt{n-r} + t}{\sqrt{n-r-0.5}} \right] \geq 1 - 1.5\mathrm{e}^{-t^2/2}$$

We may insert these probabilistic bounds into Theorem B.2. Applying the statement then ensures that the reaper algorithm (B.1) is capable of perfectly removing random white noise corruptions.

Theorem B.5 (Theorem 7.3, restatement). *Suppose we observe the matrix $\boldsymbol{B} = \begin{bmatrix} \boldsymbol{L}_0 & \boldsymbol{\Omega}_0 \end{bmatrix} \boldsymbol{\Pi} \in \mathbb{R}^{n \times M}$ where $M = m + m'$ and*

(1) The sign matrix $\boldsymbol{S} \in \{\pm 1\}^{n \times r}$ is Schur independent, with permeance $\nu(\boldsymbol{S})$;
(2) The data matrix $\boldsymbol{L}_0 \sim \mathrm{GLM}(\boldsymbol{S}, m)$ is drawn from the Gaussian loadings model;
(3) The noise matrix $\boldsymbol{\Omega}_0 \in \mathbb{R}^{n \times m'}$ is standard normal; and
(4) The permutation matrix $\boldsymbol{\Pi} \in \mathbb{R}^{M \times M}$ is arbitrary.

Assume that the parameters satisfy

$$\max\left\{1, \frac{m'}{n}\right\} \lesssim \sqrt{\nu(\boldsymbol{S})} \cdot \frac{m}{r}. \tag{B.4}$$

Then, with high probability over the randomness in the model, the orthogonal projector onto $\mathsf{L}_0 = \mathrm{range}(\boldsymbol{L}_0)$ is the unique solution to the optimization problem (B.1).

Proof. Choose $t \propto \min\{m, m'\}$ and invoke Lemmas B.3–B.4 and Rel. (B.3) to ensure

$$P(\mathsf{L}_0, \boldsymbol{L}_0) \gtrsim m\sqrt{\frac{n\nu(\boldsymbol{S})}{r}} \quad \text{and} \quad \hat{S}\big(\mathsf{L}_0^\perp, \boldsymbol{\Omega}_0\big) \|\boldsymbol{\Omega}_0\|_{S_\infty} \lesssim \max\left\{\sqrt{n}, \frac{m'}{\sqrt{n}}\right\}$$

with high probability each. Condition on these bounds to be valid. Then, Rel. (B.4) implies the deterministic recovery condition from Theorem B.2. Applying this statement ensures correct recovery of P_0 via the reaper algorithm. This argument is valid, irrespective of the specific choice of perturbation Π. $\qquad\square$

B.3. **Proof of Lemma B.3.** We use Mendelson's small ball method to derive a strong probabilistic lower bound on the permeance statistics associated with the Gaussian loadings model. The following variant of this versatile technique can be extracted from Tropp's proof [Tro15].

Theorem B.6 (Mendelson's small ball method). *Fix a set* $\mathsf{E} \subset \mathbb{R}^n$ *and let* $\boldsymbol{\varphi}_1,\ldots,\boldsymbol{\varphi}_m$ *be independent copies of a random vector* $\boldsymbol{\varphi} \in \mathbb{R}^n$. *Fix* $\xi > 0$ *and define*

$$Q_\xi\left(\mathsf{E};\boldsymbol{\varphi}\right) = \inf_{\boldsymbol{u}\in\mathsf{E}} \Pr\left[\left|\langle \boldsymbol{u},\boldsymbol{\varphi}\rangle\right| \ge \xi\right] \qquad\qquad \textit{(marginal tail function)},$$

$$W_m\left(\mathsf{E},\boldsymbol{\varphi}\right) = \mathbb{E}\sup_{\boldsymbol{u}\in\mathsf{E}}\langle \boldsymbol{u},\boldsymbol{h}\rangle \quad \textit{with} \quad \boldsymbol{h} = \frac{1}{\sqrt{m}}\sum_{i=1}^m \varepsilon_i\boldsymbol{\varphi}_i \qquad \textit{(empirical width)},$$

where $\varepsilon_1,\ldots,\varepsilon_m \in \{\pm 1\}$ *are independent Rademacher*[2] *random variables. Then, for any* $t > 0$,

$$\Pr\left[\tfrac{1}{\sqrt{m}}\inf_{\boldsymbol{u}\in\mathsf{E}}\sum_{i=1}^m \left|\langle \boldsymbol{u},\boldsymbol{\varphi}_i\rangle\right| \ge \xi\sqrt{m}Q_{2\xi}(\mathsf{E};\boldsymbol{\varphi}) - 2W_m(\mathsf{E},\varphi) - \xi t\right] \ge 1 - \mathrm{e}^{-t^2/2}.$$

The original formulation [Men15, KM15] lower-bounds the minimal ℓ_2-loss and readily follows from Lyapunov's inequality.

Lemma B.3 follows from adapting Theorem B.6 to the Gaussian loadings model. Two technical lemmas are instrumental for controlling the empirical width and the marginal tail function.

Lemma B.7. *Fix a sign matrix* $\boldsymbol{S} \in \{\pm 1\}^{n\times r}$, *$m$ independent standard Gaussian vectors* $\boldsymbol{g}_1,\ldots,\boldsymbol{g}_m \in \mathbb{R}^r$ *and m independent Rademacher random variables* $\epsilon_1,\ldots,\epsilon_m$. *Then,*

$$\mathbb{E}\sup_{\hat{\boldsymbol{u}}\in\mathrm{ran}(\boldsymbol{S}),\|\hat{\boldsymbol{u}}\|_{\ell_2}=1}\langle \hat{\boldsymbol{u}},\ \tfrac{1}{\sqrt{m}}\sum_{i=1}^m \epsilon_i\boldsymbol{S}\boldsymbol{g}_i\rangle \le \sqrt{n}.$$

Proof. Note that the vector $\boldsymbol{h} = \frac{1}{\sqrt{m}}\sum_{i=1}^m \epsilon_i\boldsymbol{S}\boldsymbol{g}_i$ is necessarily contained in the range of the sign matrix $\boldsymbol{S}$. Combine this insight with Jensen's inequality and independence of the Rademacher random variables $(\mathbb{E}\epsilon_i\epsilon_j = \delta_{i,j})$:

$$\mathbb{E}\sup_{\boldsymbol{u}\in\mathrm{ran}(\boldsymbol{S}),\|\boldsymbol{u}\|_{\ell_2}=1}\langle \boldsymbol{u},\ \boldsymbol{h}\rangle = \mathbb{E}\|\boldsymbol{h}\|_{\ell_2} \le \left(\mathbb{E}\|\boldsymbol{h}\|_{\ell_2}^2\right)^{1/2} = \left(\frac{1}{m}\sum_{i=1}^m \mathbb{E}\langle \boldsymbol{S}\boldsymbol{g}_i,\boldsymbol{S}\boldsymbol{g}_i\rangle\right)^{1/2} = \left(\mathbb{E}\langle \boldsymbol{S}\boldsymbol{g},\boldsymbol{S}\boldsymbol{g}\rangle\right)^{1/2}.$$

Finally, we apply isotropy of standard normal vectors $(\mathbb{E}\boldsymbol{g}\boldsymbol{g}^{\mathrm{t}} = \mathbf{I})$ to conclude

$$\mathbb{E}\langle \boldsymbol{S}\boldsymbol{g},\boldsymbol{S}\boldsymbol{g}\rangle = \tfrac{1}{r}\mathbb{E}\langle \boldsymbol{g},\ \boldsymbol{S}^{\mathrm{t}}\boldsymbol{S}\boldsymbol{g}\rangle = \tfrac{1}{r}\mathrm{tr}\left(\boldsymbol{S}^{\mathrm{t}}\boldsymbol{S}\right) = n.$$

$\qquad\square$

Lemma B.8. *Fix a sign matrix* $\boldsymbol{S} \in \{\pm 1\}^{n\times r}$ *with permeance* $\nu(\boldsymbol{S})$ *and sample a standard normal vector* $\boldsymbol{g} \in \mathbb{R}^r$. *Then,*

$$\Pr\left[\left|\langle \hat{\boldsymbol{u}},\ \boldsymbol{S}\boldsymbol{g}\rangle\right| \ge \tfrac{2}{3}\sqrt{\nu(\boldsymbol{S})n}\right] \ge \tfrac{1}{2} \quad \textit{for all unit vectors} \quad \hat{\boldsymbol{u}} \in \mathrm{ran}(\boldsymbol{S}).$$

Proof. Use the reformulation of $\nu(\boldsymbol{S})$ provided in Remark B.1 to infer $\frac{2}{3}\sqrt{\nu(\boldsymbol{S})n} \le \frac{2}{3}\|\boldsymbol{S}^{\mathrm{t}}\hat{\boldsymbol{u}}\|_{\ell_2}$ for all unit vectors $\hat{\boldsymbol{u}} \in \mathrm{ran}(\boldsymbol{S})$. Combine this with rotation invariance to deduce

$$\Pr\left[\left|\langle \hat{\boldsymbol{u}},\ \boldsymbol{S}\boldsymbol{g}\rangle\right| \ge 2\sqrt{\tfrac{\nu(\boldsymbol{S})n}{9r}}\right] \ge \Pr\left[\left|\langle \boldsymbol{S}^{\mathrm{t}}\hat{\boldsymbol{u}},\boldsymbol{g}\rangle\right| \ge \tfrac{2}{3}\|\boldsymbol{S}^{\mathrm{t}}\hat{\boldsymbol{u}}\|_{\ell_2}\right] = \Pr\left[\left|\langle \mathbf{e}_1,\ \boldsymbol{g}\rangle\right| \ge \tfrac{2}{3}\right] \quad \text{for all} \quad \hat{\boldsymbol{u}} \in \mathrm{ran}(\boldsymbol{S}).$$

Well-known results state that the median of the half-normal random variable $|\langle \mathbf{e}_1,\ \boldsymbol{g}\rangle|$ is $\sqrt{2}\mathrm{erf}^{-1}(1/2) > \frac{2}{3}$. Apply the defining property of the median to deduce the claim. $\qquad\square$

---

[2]A Rademacher random variable takes the values $+1$ and -1 with equal probability.

Proof of Lemma B.3. Fix a sign matrix $S \in \{\pm 1\}^{n \times r}$, a natural number $m \geq r$ and suppose that $L_0 = \begin{bmatrix} l_1 & \dots & l_m \end{bmatrix} \sim$ GLM(S, m) is sampled from the Gaussian loadings model. By assumption, each column $l_i \in \mathbb{R}^n$ $(1 \leq i \leq m)$ is an independent copy of the random vector

$$l = \frac{1}{\sqrt{r}} S g \quad \text{where} \quad g \in \mathbb{R}^r \quad \text{is standard normal.}$$

Set $\mathsf{L}_0 = \operatorname{ran}(L_0) \subset \mathbb{R}^n$ and denote its intersection with the unit sphere by $\mathsf{E} = \{\hat{u} \in \mathsf{L}_0 : \|\hat{u}\|_{\ell_2} = 1\}$. This allows us to rewrite the permeance statistics (B.2) as

$$P(\mathsf{L}_0, L_0) = \inf_{\hat{u} \in \mathsf{E}} \sum_{i=1}^m |\langle \hat{u}, l_i \rangle|.$$

Theorem B.6 lower-bounds expressions of precisely this form. Note that $\mathsf{L}_0 \subseteq \operatorname{ran}(S)$, regardless of the value of m. Lemma B.7 then asserts

$$W_m(\mathsf{E}, l) \leq \mathbb{E} \sup_{\hat{u} \in \operatorname{ran}(S), \|\hat{u}\|_{\ell_2} = 1} \langle \hat{u}, \tfrac{1}{\sqrt{m}} \sum_{i=1}^m \epsilon_i S g_i \rangle \leq \sqrt{n},$$

while Lemma B.8 ensures that for $\xi_0 = \sqrt{\frac{\nu(S)n}{9r}}$, the marginal tail function obeys

$$Q_{2\xi_0}(\mathsf{E}, l) = \inf_{\hat{u} \in \mathsf{E}} \Pr\left[|\langle \hat{u}, l \rangle| \geq \frac{2}{3\sqrt{r}} \sqrt{\nu(S)n}\right] \geq \inf_{\hat{u} \in \operatorname{ran}(S), \|\hat{u}\|_{\ell_2} = 1} \Pr\left[|\langle \hat{u}, S g \rangle| \geq \frac{2}{3} \sqrt{\nu(S)n}\right] \geq \frac{1}{2}.$$

Insert these bounds into the assertion of Theorem B.6 to complete the proof. $\qquad\square$

REFERENCES

[Bac13]   F. R. Bach. Convex relaxations of structured matrix factorizations. Available at http://arXiv.org/abs/1309.3117, 2013.

[BE16]   M. Basbug and B. Engelhardt. Hierarchical compound poisson factorization. In M. F. Balcan and K. Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1795–1803, New York, New York, USA, 20–22 Jun 2016. PMLR.

[Bha97]   R. Bhatia. *Matrix analysis*, volume 169 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1997.

[Bha07]   R. Bhatia. *Positive definite matrices*. Princeton Series in Applied Mathematics. Princeton University Press, Princeton, NJ, 2007.

[Bru17]   J. J. Bruer. *Recovering structured low-rank operators using nuclear norms*. PhD thesis, Caltech, Pasadena, 2017.

[CDS02]   M. Collins, S. Dasgupta, and R. E. Schapire. A generalization of principal component analysis to the exponential family. In *Adv. Neural Information Processing Systems*, 2002.

[CLMW11] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, 2011.

[Com94]   P. Comon. Independent component analysis, a new concept? *Signal Process*, 36(3):287 – 314, 1994.

[CSPW11] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM J. Optimiz.*, 21(2):572–596, 2011.

[DS01]   K. R. Davidson and S. J. Szarek. Chapter 8 - local operator theory, random matrices and banach spaces. In W. Johnson and J. Lindenstrauss, editors, *Handbook of the Geometry of Banach Spaces*, volume 1 of *Handbook of the Geometry of Banach Spaces*, pages 317 – 366. Elsevier Science B.V., 2001.

[Hig08]   N. J. Higham. *Functions of matrices*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2008. Theory and computation.

[HR09]   P. J. Huber and E. M. Ronchetti. *Robust statistics*. Wiley Series in Probability and Statistics. John Wiley & Sons, Inc., Hoboken, NJ, second edition, 2009.

[HV19]   B. D. Haeffele and R. Vidal. Structured low-rank matrix factorization: Global optimality, algorithms, and applications. *IEEE T Pattern Anal*, pages 1–1, 2019.

[Jag11]   M. Jaggi. *Sparse convex optimization methods for machine learning*. PhD thesis, ETH Zürich, 2011.

[Jol02]   I. T. Jolliffe. *Principal component analysis*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2002.

[KM15]   V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *Int. Math. Res. Not. IMRN*, (23):12991–13008, 2015.

[KT19]   R. Kueng and J. Tropp. Binary component decomposition Part I: The positive-semidefinite case. *preprint*, 2019.

[LMTZ15] G. Lerman, M. B. McCoy, J. A. Tropp, and T. Zhang. Robust computation of linear models by convex relaxation. *Found. Comput. Math.*, 15(2):363–410, 2015.

[LP96] M. Laurent and S. Poljak. On the facial structure of the set of correlation matrices. *SIAM J. Matrix Anal. Appl.*, 17(3):530–547, 1996.

[Men15] S. Mendelson. Learning without concentration. *J. ACM*, 62(3):Art. 21, 25, 2015.

[Mir60] L. Mirsky. Symmetric gauge functions and unitarily invariant norms. *Quart. J. Math. Oxford Ser. (2)*, 11:50–59, 1960.

[OF96] B. A. Olshausen and D. J. Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607–609, 1996.

[PT94] P. Paatero and U. Tapper. Positive matrix factorization: A non-negative factor model with optimal utilization of error estimates of data values. *Environmetrics*, 5(2):111–126, 1994.

[Sre04] N. Srebro. *Learning with matrix factorizations*. ProQuest LLC, Ann Arbor, MI, 2004. Thesis (Ph.D.)–Massachusetts Institute of Technology.

[TB99] M. E. Tipping and C. M. Bishop. Probabilistic principal component analysis. *J. R. Stat. Soc. Ser. B Stat. Methodol.*, 61(3):611–622, 1999.

[Tik18] K. Tikhomirov. Singularity of random bernoulli matrices. *arXiv preprint arXiv:1812.09016*, 2018.

[Tro04] J. A. Tropp. *Topics in sparse approximation*. PhD thesis, The University of Texas at Austin, 2004.

[Tro15] J. A. Tropp. *Convex Recovery of a Structured Signal from Independent Random Linear Measurements*, pages 67–101. Springer International Publishing, Cham, 2015.

[Tro18] J. A. Tropp. Simplicial faces of the set of correlation matrices. *Discrete Comput. Geom.*, 60(2):512–529, 2018.

[Ude15] M. Udell. *Generalized low rank models*. PhD thesis, Stanford University, 2015.

[Wit10] D. M. Witten. *A penalized matrix decomposition and its applications*. PhD thesis, Stanford University, 2010.

[ZHT06] H. Zou, T. Hastie, and R. Tibshirani. Sparse principal component analysis. *J. Comput. Graph. Statist.*, 15(2):265–286, 2006.