Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

SIAM J. MATRIX ANAL. APPL.
Vol. 46, No. 3, pp. 2140–2170
© 2025 Society for Industrial and Applied Mathematics

ON THE ROBUSTNESS OF THE SUCCESSIVE PROJECTION ALGORITHM*

GIOVANNI BARBARINO† AND NICOLAS GILLIS†

Abstract. The successive projection algorithm (SPA) is a workhorse algorithm to learn the r vertices of the convex hull of a set of $(r-1)$-dimensional data points, a.k.a. a latent simplex, which has numerous applications in data science. In this paper, we revisit the robustness to noise of SPA and several of its variants. In particular, when $r \geq 3$, we prove the tightness of the existing error bounds for SPA and for two more robust preconditioned variants of SPA. We also provide significantly improved error bounds for SPA, by a factor proportional to the conditioning of the r vertices, in two special cases: for the first extracted vertex and when $r \leq 2$. We then provide further improvements for the error bounds of a translated version of SPA proposed by Arora et al. [*Proceedings of the International Conference on Machine Learning*, 2013, pp. 280–288] in two special cases: for the first two extracted vertices and when $r \leq 3$. Finally, we propose a new more robust variant of SPA that first shifts and lifts the data points in order to minimize the conditioning of the problem. We illustrate our results on synthetic data.

Key words. simplex-structured matrix factorization, successive projection algorithm, nonnegative matrix factorization, latent simplex, low-rank matrix approximations

MSC codes. 15A23, 65F50, 65F25, 68Q25, 65D18

DOI. 10.1137/24M171293X

Code and Data Available

See reproducibility of computational results at end of the article.

1. Introduction.

The problem of finding a latent simplex via the observation of noisy data points generated within that simplex is a fundamental problem in signal processing, data analysis, and machine learning. It finds applications in chemometrics [3, 33], hyperspectral imaging [26], audio source separation [36], topic modeling [4, 5], and community detection [27, 21, 2], to cite a few. The problem is posed as follows: we observe samples

$$x_j = Wh_j + N(:,j) \in \mathbb{R}^m \text{ for } j = 1, 2, \ldots, n,$$

where the following hold:

- The columns of $W \in \mathbb{R}^{m \times r}$ are r vertices, to be found, $\{w_k \in \mathbb{R}^m\}_{k=1}^r$.
- The weights $h_j \in \mathbb{R}^r$ are nonnegative and sum to one, that is, they belong to the probability simplex, $h_j \in \Delta = \{y \in \mathbb{R}^r \mid y \geq 0, \sum_{i=1}^r y_i = 1\}$ for all j. We will denote $H = [h_1, h_2, \ldots, h_n] \in \mathbb{R}^{r \times n}$.
- The vector $N(:,j) \in \mathbb{R}^m$ models the noise and model misfit.

*Received by the editors November 25, 2024; accepted for publication (in revised form) June 20, 2025; published electronically September 19, 2025.
https://doi.org/10.1137/24M171293X
Funding: The authors acknowledge support by the European Union (ERC consolidator, eLinoR, 101085607).
†Departement of Mathematics and Operational Research, University of Mons, 7000 Mons, Belgium (giovanni.barbarino@umons.ac.be, nicolas.gillis@umons.ac.be).

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Algorithm 1. Successive projection algorithm [3].

Input: The matrix $X = WH + N$, where $H \geq 0$, $H^\top e \leq e$, H is separable, W is full column rank, the number r of columns to extract.

Output: Set of r indices $\mathcal{J}$ such that $\tilde{X}(:, \mathcal{J}) \approx W$ (up to permutation).

Let $R = X$, $\mathcal{J} = \{\}$.
for $k = 1, \ldots r$ **do**
  $p \in \operatorname{argmax}_j \|R(:, j)\|$. *% Selection step*
  $\mathcal{J} \leftarrow \mathcal{J} \cup \{p\}$.
  $R \leftarrow (I - \frac{R(:,p)R(:,p)^\top}{\|R(:,p)\|^2})R$. *% Projection step*
end for

Hence the samples, x_j's, are contained in a latent simplex,[1] $\operatorname{conv}(W) = \{x \mid x = Wh, h \in \Delta^r\}$, up to the noise $N(:, j)$, and the goal is to learn W given $X = WH + N$. This problem is sometimes referred to as "learning a latent simplex," [5] "simplex-structured matrix factorization" (SSMF) [1], or "probabilistic simplex component analysis" [38]; we will use SSMF in this paper. When $W \geq 0$, this is a nonnegative matrix factorization (NMF) problem [24].

Many variants of SSMF have been studied, making various assumptions on W, H, and N. An assumption that has proved particularly powerful is *separability*, which requires that for each vertex (that is, each column of W), there exists at least one data point close to that vertex. Mathematically: for all $k = 1, 2, \ldots, r$, there exists j such that $X(:, j) = W(:, k) + N(:, j)$. This is equivalent to requiring that H contains the identity matrix as a submatrix since the previous equation implies $H(:, j) = e_k$, where e_k the kth unit vector. If H satisfies this condition, it is said to be a separable matrix. The separability assumption is also known as the pure-pixel assumption in hyperspectral imaging [26] or the anchor-word assumption in topic modeling [4]. In this paper, we will refer to the SSMF problem under the separability assumption as *separable SSMF.*

Countless algorithms have been developed for separable SSMF; see, e.g., [15, Chapter 7] and the references therein. Among them, a workhorse algorithm is the successive projection algorithm (SPA) [3], which is described in Algorithm 1. Geometrically, when $\operatorname{rk}(W) = r$, $\operatorname{conv}(W)$ is an $(r-1)$-dimensional simplex with r vertices embedded in a higher m-dimensional space. The first step of SPA extracts the point with the largest ℓ_2 norm. For any polytope (such as a simplex), such a point is always a vertex. Then SPA projects all data points on the orthogonal complement of that extracted point (so that this point is projected onto zero), reducing the dimension of the simplex by one, and continuing the process $(r-1)$ times.

Remark 1 (implementation of SPA). SPA should not be implemented as in Algorithm 1; this would be particularly ineffective for large sparse matrices, as the residual would become quickly dense (especially if the first extracted columns is dense). It can be implemented in $\mathcal{O}(r \operatorname{nnz}(X))$ operations (essentially, r matrix-vector products), where $\operatorname{nnz}(X)$ is the number of nonzero entries of X, using the fact that

[1] A simplex is a d-dimensional poltyope with $d+1$ vertices, e.g., the probability simplex. For $\operatorname{conv}(W) \subset \mathbb{R}^m$ to be a simplex embedded in an m-dimensional space, the affine hull of the columns of W, that is, $\{Wx \mid x^\top e = 1\}$, must have dimension $r-1$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$\|(I - uu^\top)y\|^2 = \|y\|^2 - (u^\top y)^2$, where u has unit norm and y is any vector [19]. See https://gitlab.com/ngillis/robustSPA for a MATLAB implementation.

SPA has a long history, as the index set of columns it extracts corresponds to that of QR with column pivoting [8]. However, it was introduced in the context of separable SSMF by Araújo et al. [3] and rediscovered many times under different names; see [15, p. 229] for an historical overview.

Besides being simple and computationally very cheap, SPA was one of the first algorithms for separable SSMF to be proved to be robust to noise [19]; see the next section for more details. Let us recall this result.

THEOREM 1 ([19, Theorem 3]). *Let $X = WH + N \in \mathbb{R}^{m\times n}$, where $H \in \mathbb{R}_+^{r\times n}$ is a separable matrix satisfying $H^\top e \le e$ with e the vector of all ones of appropriate dimension, and let*

$$\text{(Bound on the noise)} \qquad \varepsilon = \max_j \|N(:,j)\| \le \mathcal{O}\left(\frac{\sigma_r(W)}{\sqrt{r}\mathcal{K}^2(W)}\right),$$

where $\sigma_r(W)$ is the rth singular value of W, and $\mathcal{K}(W) = \frac{K(W)}{\sigma_r(W)}$ with $K(W) = \max_k \|W(:,k)\|$ is a measure of the conditioning[2] of W. Then SPA (Algorithm 1) returns a set of indices $\mathcal{J}$ such that

$$\text{(Error bound)} \qquad \max_{1\le k\le r} \min_j \|W(:,k) - X(:,\mathcal{J}_j)\| \le \mathcal{O}\left(\varepsilon\mathcal{K}^2(W)\right).$$

Note that it is not required that H is column stochastic, that is, $H^\top e = e$, but only that $H^\top e \le e$; we will discuss this in section 2.

Outline and contribution. In this paper, we revisit the robustness to noise of SPA and some of its variants and provide a much more complete picture than that of Theorem 1. After summarizing the current knowledge about the robustness to noise of SPA in section 2, our contributions are as follows.

- Section 3. We reanalyze the selection step of SPA, showing that the error bound of SPA can be improved, namely, in $\mathcal{O}(\varepsilon\mathcal{K}(W))$ instead of $\mathcal{O}(\varepsilon\mathcal{K}^2(W))$ for the first step of SPA (Theorem 3). We then show that for a rank-2 separable SSMF problem, SPA achieves the error bound in $\mathcal{O}(\varepsilon\mathcal{K}(W))$ with a bound on the noise given by $\varepsilon \le \mathcal{O}(\frac{\sigma_r(W)}{\mathcal{K}(W)})$, a significant improvement compared to SPA, of a factor $\mathcal{K}(W)$ for both terms (Theorem 4).
- Section 4. For separable SSMF, the first projection step of SPA can be replaced by a translation, as is done in [4], which we refer to as translated SPA (T-SPA). Adapting the results from section 3, we show that the first two steps of T-SPA have an error bound in $\mathcal{O}(\varepsilon\mathcal{K}(W))$ (Theorem 5). Moreover, for a rank-3 separable SSMF problem, T-SPA achieves the error bound in $\mathcal{O}(\varepsilon\mathcal{K}(W))$ with a bound on the noise given by $\varepsilon \le \mathcal{O}(\frac{\sigma_r(W)}{\mathcal{K}(W)})$, the same improvement as for SPA in section 3 but for $r = 3$ (Theorem 6).
- Section 5. For $r \ge 3$, we show that the bound of Theorem 1 from [19] with error bound in $\mathcal{O}(\varepsilon\mathcal{K}^2(W))$ is tight (Theorem 7). We do this by providing a family of data matrices that achieve this worst-case bound already at the second step of SPA. This answers an open question about the tightness of the SPA error bound of Theorem 1 [15, p. 231].

[2]Note that $\mathcal{K}(W)$ is smaller than the condition number of W, $\frac{\sigma_1(W)}{\sigma_r(W)}$. In fact, recall that $\sigma_1(W) = \|W\| \ge K(W) \ge \sigma_1(W)/\sqrt{r}$, where r is the number of columns in W.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

- Section 6. SPA can be preconditioned in various ways to obtain error bounds in $\mathcal{O}(\varepsilon\mathcal{K}(W))$ [20, 18] (see section 2.2 for more details). We show that these error bounds are tight for two important preconditioners: SPA itself [18] (Theorem 9) and minimum-volume ellipsoid (MVE) [20] (Theorem 11).
- Section 7. We discuss a way to preprocess the data, by translating and lifting the data points in a higher dimensional space of dimension $m+1$, allowing us to improve robustness of SPA be reducing the conditioning of W.
- Section 8. We provide numerical experiments to compare the various SPA variants and how they compare on synthetic data, allowing us to illustrate our theoretical findings. In particular, our experiments show that our proposed SPA variant outperforms the others in adversarial settings.

Notation. Given a vector $x \in \mathbb{R}^m$, we denote $\|x\|$ its ℓ_2 norm. Given a matrix $X \in \mathbb{R}^{m\times n}$, we denote $X^\top$ its transpose, $X(:,j)$ its jth column, $X(i,:)$ its ith row, $X(i,j)$ its entry as position (i,j), $\|X\| = \sigma_{\max}(X)$ its ℓ_2 norm which is equal to its largest singular value, $\|X\|_F^2 = \sum_{i,j} X(i,j)^2$ its squared Frobenius norm, $\sigma_r(X)$ its rth singular value, $\sigma_{\min}(X) = \sigma_{\min(m,n)}(X)$ its smallest singular value, $\mathrm{rk}(X)$ its rank, $K(X) = \max_j\{\|X(:,j)\|\}$, and $\mathcal{K}(X) = K(X)/\sigma_{\min}(X)$. For $k > \min\{m,n\}$, by convention, $\sigma_k(X) = 0$. Note that

$$\sigma_n(X) \le \min_i\{\|X(:,i)\|\} \le K(X) \le \|X\|,$$
$$\sigma_n(X) \le \|P_{X(:,i)}(X(:,j))\| \le \|X(:,i) - X(:,j)\|,$$

and $\sigma_n(X)\sqrt{2} \le \|X(:,i) - X(:,j)\|$ for every $i \ne j$, where $P_u(v)$ is the projection of v on the subspace orthogonal to u. We denote e_k the kth unit vector, e the vector of all ones of appropriate dimension, and I_r the identity matrix of dimension r. The set $\mathbb{R}_+^{m\times n}$ denotes the m-by-n componentwise nonnegative matrices. A matrix $H \in \mathbb{R}^{r\times n}$ has stochastic columns if $H \ge 0$ and $H^\top e = e$. It is separable if $H(:,\mathcal{J}) = I_r$ for some index set $\mathcal{J}$ of size r.

Given $W \in \mathbb{R}^{m\times r}$, we may also denote its columns by $\{w_j\}_{j=1}^r$. The convex hull generated by the columns of W is denoted $\mathrm{conv}(W) = \mathrm{conv}(w_1, w_2, \ldots, w_r) = \{Wh \mid e^\top h = 1, h \ge 0\}$.

The notation $A \succ 0$ means that A is a square symmetric positive definite matrix.

In the results, we use the notation $\cdots \le \mathcal{O}(f(\mathbf{x}))$ to indicate that a quantity is bounded by $f(\mathbf{x})$ up to an absolute multiplicative constant C that does not depend on any of the parameters or variables.

2. SPA: What do we know?

As described in the previous section, SPA is fast and robust to noise. However, it has several drawbacks, including the following ones:

1. It can only extract as many as $\mathrm{rk}(W)$ vertices. For example, if the columns of W live in an m-dimensional space with $\mathrm{rk}(W) \le m < r$, e.g., a triangle in the plane ($m = \mathrm{rk}(W) = 2$, $r = 3$), SPA can only extract $\mathrm{rk}(W)$ of the r vertices.
2. The error bounds in Theorem 1 are rather poor, depending on the squared of the condition number of W. Moreover, it is not known whether these bounds are tight [15, p. 231].
3. SPA is sensitive to outliers.

Researchers have proposed many improvements and modifications of SPA to alleviate some of these drawbacks. Let us highlight some of these results; in particular the ones we will discuss in this paper.

2.1. Drawback 1: Rank-deficient case.

How can we adapt SPA to handle the case when $\mathrm{rk}(W) < r$? Assuming that H is column stochastic, that is, $H^\top e = e$,

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Algorithm 2. Translated SPA [4].

Input: The matrix $X = WH + N$, where $H \geq 0$, $H^\top e = e$, H is separable, W is full column rank, the number r of columns to extract.

Output: Set of r indices $\mathcal{J}$ such that $X(:,\mathcal{J}) \approx W$ (up to permutation).

Let $R = X$, $\mathcal{J} = \{\}$.
for $k = 1, \ldots r$ **do**
  $p \in \operatorname{argmax}_j \|R(:,j)\|$. *% Selection step*
  $\mathcal{J} = \mathcal{J} \cup \{p\}$.
  if $k = 1$ **then**
    $R \leftarrow R - X(:,p)\, e^\top$. *% Translation step*
  else
    $R \leftarrow (I - \frac{R(:,p)R(:,p)^\top}{\|R(:,p)\|^2})R$. *% Projection step*
  end if
end for

Arora et al. [4] replace the first projection step of SPA by a translation step; see Algorithm 2. Note that Arora et al. [4] did not present their algorithm exactly in these terms (they worked with affine spaces), but Algorithm 2 is equivalent.

Remark 2 (implementation of T-SPA). As for SPA (see Remark 1), T-SPA can be implemented in $\mathcal{O}(r\,\mathrm{nnz}(X))$ operations (essentially, r matrix-vector products). See https://gitlab.com/ngillis/robustSPA for a MATLAB implementation.

Replacing the projection step by a translation at the first iteration has two advantages:

- It allows T-SPA to extract $\mathrm{rk}(W) + 1$ vertices. The simplest examples are triangles in the plane, e.g.,
$$W = \begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \end{pmatrix},$$
for which SPA can only extract 2 out of the 3 vertices, while T-SPA can extract all of them.
- Arora et al. [4] proved an error bound in $\mathcal{O}(\varepsilon \mathcal{K}^2(W_t))$ for T-SPA, where $W_t \in \mathbb{R}^{m \times (r-1)}$ is the translated version of W where the zero column is discarded, that is,
$$W_t = [w_1 - w_p, \ldots, w_{p-1} - w_p, w_{p+1} - w_p, \ldots, w_r - w_p],$$
where w_p is the first extracted column by T-SPA, that is, the column of W with the largest ℓ_2 norm. Note that this bound essentially follows that of [19] after translation. Interestingly, and this has never been pointed out in the literature as far as we know, the matrix W_t can be significantly better conditioned than W, even when $\mathrm{rk}(W) = r$. For example, let
$$W = \begin{pmatrix} -1 & 1 & 0 \\ -1 & 0 & 1 \\ \delta & \delta & \delta \end{pmatrix} \text{ for } \delta \ll 1, \text{ with } W_t = \begin{pmatrix} 2 & 1 \\ 1 & 2 \\ 0 & 0 \end{pmatrix},$$
where $\mathcal{K}(W) = \mathcal{O}(1/\delta)$ while $\mathcal{K}(W_t) = \sqrt{5}$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Note that we will always have $\mathcal{K}(W_t) \leq 2\mathcal{K}(W)$. In fact, $K(W_t) \leq 2K(W)$ since $\|w_i - w_j\| \leq \|w_i\| + \|w_j\|$, and $\sigma_{r-1}(W_t) \geq \sigma_r(W)$ since W_t is a rank-one correction of W (by the interlacing singular value theorem).

Another possibility to extract $\text{rk}(W) + 1$ vertices when the affine hull of the columns of W contains the origin is to replace X by $\binom{X}{ce^\top}$, where $c > 0$ is a constant [27, 21]. This amounts to lifting the data points in one dimension higher and leads to an equivalent SSMF problem: for $H^\top e = e$,

$$X = WH \iff \begin{pmatrix} X \\ ce^\top \end{pmatrix} = \begin{pmatrix} W \\ ce^\top \end{pmatrix} H. \tag{2.1}$$

This prevents the affine hull of $\binom{W}{ce^\top}$ from containing the origin. We will discuss in section 7 how to perform this lifting (and combining it with a translation) in order to minimize the conditioning of $\binom{W}{ce^\top}$ and hence improve the performance of SPA.

The two methods described above can only extract the vertices of a simplex, that is, of an $(r-1)$-dimensional poltyope with r vertices. In order to extract more vertices, one can take the nonnegativity of H into account in the projection step, leading to the successive nonnegative projection algorithm (SNPA). SNPA can extract all the vertices of $\text{conv}(W)$, regardless of the dimensions; see [13] for more details. However, when $\text{rk}(W) = r$, SPA and SNPA have the same error bounds, while SPA is significantly faster. Other approaches that allow one to extract more than $\text{rk}(W) + 1$ vertices are based on convex formulations [11, 10, 34, 12, 17, 29]. We will not discuss SNPA and the convex relaxations further in this paper, nor the tightness of their error bounds, which is a topic of further research.

2.2. Drawback 2: Sensitivity to noise. A standard preprocessing step in the literature to improve the robustness to noise of SPA is to replace X with its best rank-r approximation using the truncated singular value decomposition (SVD); see, e.g., [26, 22]. It is also possible to precondition the input matrix to improve the performance of SPA in the presence of noise. The intuition behind preconditioning is to try to estimate the left inverse of W, $W^\dagger$, from the data, then multiply X by $W^\dagger$, and apply SPA on $W^\dagger X = H + W^\dagger N$. In the preconditioned problem, W becomes the identity which is perfectly conditioned. In this paper, we will focus on two preconditionings because they offer the most robust variants of SPA:

- MVE [20]: find the r-dimensional MVE centered at the origin, $\{x \mid x^\top Ax \leq 1\}$, where $A \succ 0$, that contains all data point. The MVE can be computed via semidefinite programming. Then multiply X by the left inverse of A. This allows one to reduce the bound on the noise in SPA to $\varepsilon \leq \mathcal{O}(\frac{\sigma_r(W)}{r\sqrt{r}})$, with a factor of improvement of $\frac{\mathcal{K}^2(W)}{r}$, and reduce the error bound to $\mathcal{O}(\varepsilon\mathcal{K}(W))$, with a factor of improvement of $\mathcal{K}(W)$. This approach is referred to as MVE-SPA. Note that a similar approach was proposed in [28].
  Note that MVE-SPA needs to work with an r-by-n input matrix, and hence a dimensionality reduction preprocessing is necessary, such as the truncated SVD.
- SPA [18]: Estimate $W^\dagger$ by the left inverse of the SPA solution, $X(:, \mathcal{J})$. This strategy might appear a bit odd, preconditioning SPA with itself. It has the same bound on the noise as SPA but allows one to reduce the error bound to $\mathcal{O}(\varepsilon\mathcal{K}(W))$, while providing significantly better results in numerical experiments [18] (see also section 8). We will denote SPA preconditioned with SPA as SPA^2.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

We will show in section 6 that the error bounds for these two preconditionings are also tight.

Another way to improve the error bound of SPA to $\mathcal{O}(\varepsilon\mathcal{K}(W))$ (but not the bound on the noise) was proposed in [4]: after a first set of r indices $\mathcal{J}$ is extracted by SPA, each element $j \in \mathcal{J}$ is reviewed again and replaced with the index p that maximizes the volume of $\text{conv}(X(:,\mathcal{J}\backslash\{j\} \cup \{p\})$. However, this algorithm, referred to as fast anchor word (FAW), typically performs worse than SPA[2] [18]; see also section 8 for numerical experiments. Moreover, the computational cost of FAW is r times that of SPA. Note that FAW uses a translation at the first step, that is, it relies on T-SPA.

Another direction of research to improve robustness to noise is to consider randomized algorithms, such as vertex component analysis [31] or randomized SPA [37]. Such algorithms generate a different solution at each run, from which one can pick the best one according to some criterion (e.g., the reconstruction error or the volume of $\text{conv}(W)$). Discussing such algorithms is out of the scope of this paper.

2.3. Drawback 3: Outliers. To improve performance in the presence of outliers, standard strategies include outlier removal (see, e.g., [22]) and robust low-rank approximations of the input matrix [9]. Another dedicated strategy for SPA is to check, after the selection step, that the extracted column of X can sufficiently reduce the error in the approximation (that is, $\min_H \|X - X(:,\mathcal{J})H\|$ is sufficiently small) [14]. A more recent strategy leverages the presence of multiple columns of X close to that of W, allowing one to estimate the columns of W as the average or median of several columns of X; see [6, 5, 30]. We will not cover these approaches in this paper.

3. Improved bounds for the first step of SPA and when $r = 2$. In this section, we provide improved bounds for SPA: for the first step in section 3.1 (Theorem 3) and when $r = 2$ in section 3.2 (Theorem 4).

Motivations. From a theoretical standpoint, understanding better the behavior of SPA, a workhorse separable NMF algorithm, is a mathematically interesting problem, even when $r = 2, 3$. From a practical standpoint, rank-2 NMF and/or SPA with $r = 2$ have been used successfully to segment and analyze medical images [25, 35], extract topics in a large corpus of documents [23], and cluster pixels in hyperspectral images [16]. In analytical chemistry, where NMF is referred to as the self-modeling curve resolution problem, using $r = 3$ is a standard setting [7, 32].

3.1. First step of SPA. The robustness proof of SPA in Theorem 1 relies on using sequentially the following theorem that studies a single step of SPA (we report the bounds using the $\mathcal{O}$ notation to simplify the presentation[3]).

THEOREM 2 ([19, Theorem 2]). *Let $W \in \mathbb{R}^{m\times k}$, and $Q \in \mathbb{R}^{m\times(r-k)}$ and $X = [W\ Q]H + N \in \mathbb{R}^{m\times n}$ with $k \le r \le \min\{m,n\}$, where W is full rank, $K(N) < \varepsilon \le \mathcal{O}(\frac{\sigma_k(W)}{\mathcal{K}(W)})$, and $H \in \mathbb{R}_+^{r\times n}$ is separable with $H^\top e \le e$. If $K(Q) \le \mathcal{O}(\mathcal{K}(W)\varepsilon)$, then the column of X with the largest ℓ_2 norm, denoted x, satisfies*

$$\min_{1\le i\le r} \|x - w_i\| \le \mathcal{O}\left(\mathcal{K}^2(W)\varepsilon\right).$$

In this section, we improve the bound of Theorem 2; see Theorem 3 below, which is very similar to Theorem 2 but improves the error bound, namely, from $\mathcal{O}(\varepsilon\mathcal{K}^2(W))$ to $\mathcal{O}(\varepsilon\mathcal{K}(W))$. Note that it does not improve [19, Theorem 2] in terms of the bound on the noise allowed by SPA, which was already of the order of $\varepsilon \le \mathcal{O}(\frac{\sigma_r(W)}{\mathcal{K}(W)})$.

[3]In [19], the authors consider any strongly convex and L-smooth function f to be maximized, not only the ℓ_2 norm. We focus in this paper on the ℓ_2 norm, which is the most used variant.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

THEOREM 3. *Let $W \in \mathbb{R}^{m\times k}$, $Q \in \mathbb{R}^{m\times(r-k)}$, and $X = [W\ Q]H + N \in \mathbb{R}^{m\times n}$ with $k \le \min\{m,r\}$, where W is full rank, $K(N) < \varepsilon \le \sigma_k(W)\mathcal{K}(W)^{-1}/8$, and $H \in \mathbb{R}_+^{r\times n}$ is separable and has stochastic columns. If $K(Q) \le K(W)/2$, then for the column of X with largest ℓ_2 norm, denoted $x = X(:,p) = [W,Q]h_p + N(:,p)$ for some p, there exists a column w_i of W such that $\|x - w_i\| \le 33\mathcal{K}(W)\varepsilon$ and $h_p(i) \ge 1/2$.*

Notice that under these hypotheses, $K(W)/2 \ge \sigma_k(W)/2 \ge 4\mathcal{K}(W)\varepsilon$, so the upper bound on $K(Q)$ is consistent with the one in Theorem 2 and possibly less restrictive.

To prove Theorem 3, we use a similar proof strategy as that of [22, Theorem 1], which is based on the geometry of the problem, as opposed to matrix norm inequalities used in [19], as explained in [22]. However, [22, Theorem 1] still has a factor proportional to $\mathcal{K}^2(W)$ in the error bound, and hence its bound is asymptotically similar to that of Theorem 2. Moreover, [22, Theorem 1] claims that it can replace $\sigma_r(W)$ in Theorem 1 by $\sigma_{r-1}(W)$ in the bound on the noise and in the error bound. However, this is not possible, as exemplified by the following counterexample. Suppose in fact that $W = \left(\begin{smallmatrix} 0 & 0 \\ 1 & 0 \end{smallmatrix}\right)$, $r = 2$ and $X = \left(\begin{smallmatrix} 0 & \varepsilon & 0 \\ 1 & 1 & 0 \end{smallmatrix}\right)$ for a small enough $\varepsilon > 0$. Notice that $X = WH + N$ for a separable and column stochastic $H = \left(\begin{smallmatrix} 1 & 1 & 0 \\ 0 & 0 & 1 \end{smallmatrix}\right)$ and with $K(N) = \varepsilon$. Then SPA identifies the first two columns of X as the vertices of the simplex, and the error is at least 1 for any $\varepsilon > 0$; this does not coincide with the estimation $\mathcal{O}(\varepsilon K(W)^2/\sigma_{r-1}(W)^2) = \mathcal{O}(\varepsilon)$.

Before proving Theorem 3, let us state the following two lemmas.

LEMMA 1. *Let $W \in \mathbb{R}^{m\times r}$ with $r \le m$. If $\varepsilon \le \sigma_r(W)\mathcal{K}(W)^{-1}/8$, then for every couple of distinct indices i,j,*

$$\left\| \frac{w_i + w_j}{2} \right\| \le \max\{\|w_i\|, \|w_j\|\} - 2\varepsilon.$$

Proof. See Appendix A.1. ∎

LEMMA 2. *Let V be a convex polytope with vertices $\{v_0, v_1, \ldots, v_r\}$. Let C be a closed convex set containing $v_1, \ldots, v_r$ but not v_0. Let $w_i = v_0 + \alpha_i(v_i - v_0)$ where $0 < \alpha_i \le 1$ is the smallest value for which $w_i \in C$. Then the set $V \setminus C$ is contained in* $\operatorname{conv}(v_0, w_1, \ldots, w_r)$.

Proof. See Appendix A.2. ∎

We can now prove Theorem 3.

Proof of Theorem 3. The vector x is equal to $w + s$, where $w = [W\ Q]h$ and h, s are columns of H and N, respectively, so $\|x\| < \|w\| + \varepsilon$. At the same time, given $\|w_i\| = K(W)$ the longest column of W, there exists a column of X that is equal to w_i up to a perturbation of norm at most ε, due to the separability of the matrix H, so $\|x\| > K(W) - \varepsilon$. It implies that there exists a vector w in the convex hull of $[W\ Q]$ for which

$$w \in \mathscr{S} := \{y \in \operatorname{conv}([W\ Q]) \mid \|y\| > K(W) - 2\varepsilon\}.$$

Since

$$K(W) - 4\varepsilon \ge K(W) - \sigma_k(W)\mathcal{K}(W)^{-1}/2 \ge K(W)/2 \ge K(Q),$$

then all vectors in $\operatorname{conv}(Q)$ have norm at most $K(W) - 4\varepsilon$, and $\|w_i + q_j\|/2 \le K(W) - 2\varepsilon$. Call $M := [W\ Q]$ so that $m_i = w_i$ for $i \le k$ and $m_i = q_{i-k}$ for $i > k$. By the hypothesis on ε and Lemma 1, all the vectors in the polytope $\{y = [W\ Q]h \mid e^\top h = 1,$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$h \geq 0, \|h\|_\infty \leq 1/2\}$ have norm at most $K(W) - 2\varepsilon$, since its vertices are the median points between different columns of $[W\ Q]$. As a consequence,

$$\mathscr{S} = \bigcup_i \big(\mathscr{S} \cap \{y = [W\ Q]h \mid e^\top h = 1,\, h \geq 0, h_i \geq 1/2\}\big) := \bigcup_i (\mathscr{S} \cap \mathscr{V}_i) := \cup_{i \leq k} \mathscr{S}_i,$$

where the union is disjoint. Notice that the $\mathscr{S}_i$ associated with $\|m_i\| \leq K(W) - 2\varepsilon$ are empty since the vertices of $\mathscr{V}_i$ are the median points between m_i and the columns of M (itself included). In particular $\mathscr{S}_i = \emptyset$ for $i > k$ since $K(Q) \leq K(W) - 4\varepsilon$. Let now $\|w_i\| > K(W) - 2\varepsilon$. Since the ball with radius $K(W) - 2\varepsilon$ is convex and closed and contains $(w_i + m_j)/2$ for all $j \neq i$, but not w_i, we can find points $\widetilde{m}_{ij}$ of norm exactly $K(W) - 2\varepsilon$ on the segments connecting w_i with $(w_i + m_j)/2$. By convention, denote $\widetilde{m}_{ii} := w_i$. Since $\mathscr{S}_i = \mathscr{S} \cap \mathscr{V}_i = (\mathrm{conv}(M) \setminus B(K(W) - 2\varepsilon)) \cap \mathscr{V}_i = \mathscr{V}_i \setminus B(K(W) - 2\varepsilon)$, we can use Lemma 2 and find that $\mathscr{S}_i \subseteq \mathrm{conv}\{\widetilde{m}_{i1}, \ldots, \widetilde{m}_{ir}\}$. The point w is in one of the nonempty $\mathscr{S}_i$, so

$$\min_i \|x - w_i\| = \min_i \|w + s - w_i\| \leq \varepsilon + \max_{i:\mathscr{S}_i \neq \emptyset} \max_{y \in \mathscr{S}_i} \|y - w_i\| \leq \varepsilon + \max_{i:\mathscr{S}_i \neq \emptyset} \max_{j \neq i} \|\widetilde{m}_{ij} - w_i\|. \tag{3.1}$$

Let $\delta_{ij} \in (0, 1/2]$ be the value such that $\widetilde{m}_{ij} = w_i + \delta_{ij}(m_j - w_i)$ for $j \neq i$, and notice that $\|\widetilde{m}_{ij} - w_i\| = \delta_{ij}\|m_j - w_i\|$. From $\|\widetilde{m}_{ij}\| = K(W) - 2\varepsilon$, we obtain the following equation in δ_{ij}:

$$(K(W) - 2\varepsilon)^2 = \|w_i\|^2 + \delta_{ij}^2 \|m_j - w_i\|^2 + 2\delta_{ij} w_i^\top (m_j - w_i). \tag{3.2}$$

If $\|m_j\| \leq K(W) - 2\varepsilon$, then δ_{ij} will be the only solution between $[0, 1]$. Instead, if $\|m_j\| > K(W) - 2\varepsilon$, then the two solutions of the above equation will be δ_{ij} and $1 - \delta_{ji}$, where $\delta_{ij} \leq \frac{1}{2} \leq 1 - \delta_{ji}$. In this case, if we further assume $\|w_i\| \geq \|m_j\|$, we get

$$\delta_{ij} + 1 - \delta_{ji} = \frac{2w_i^\top (w_i - m_j)}{\|m_j - w_i\|^2} = \frac{2\|w_i\|^2 - 2w_i^\top m_j}{\|m_j - w_i\|^2} \geq \frac{\|w_i\|^2 + \|m_j\|^2 - 2w_i^\top m_j}{\|m_j - w_i\|^2} = 1,$$

so $\delta_{ij} \geq \delta_{ji}$ and

$$\|\widetilde{m}_{ij} - w_i\| = \delta_{ij}\|m_j - w_i\| \geq \delta_{ji}\|m_j - w_i\| = \|\widetilde{m}_{ji} - m_j\|.$$

Thus, to estimate (3.1), we just need to consider the indices i, j such that $\|w_i\| \geq \max\{\|m_j\|, K(W) - 2\varepsilon\}$. For such couples of indices, δ_{ij} will always be the smallest solution of (3.2), and $w_i^\top (w_i - m_j) \geq \|m_j - w_i\|^2/2 > 0$. From the inequality $a - \sqrt{a^2 - x} \leq x/a$ that holds for any $0 \leq x \leq a^2$ and $0 < a$,

$$\begin{aligned}
\delta_{ij} &= \frac{w_i^\top (w_i - m_j) - \sqrt{(w_i^\top (w_i - m_j))^2 - \|m_j - w_i\|^2(\|w_i\|^2 - (K(W) - 2\varepsilon)^2)}}{\|m_j - w_i\|^2} \\
&\leq \frac{\|w_i\|^2 - (K(W) - 2\varepsilon)^2}{w_i^\top (w_i - m_j)} \leq \frac{2(\|w_i\| - (K(W) - 2\varepsilon))(\|w_i\| + (K(W) - 2\varepsilon))}{\|m_j - w_i\|^2} \\
&\leq \frac{8\|w_i\|\varepsilon}{\|m_j - w_i\|^2} \leq \frac{8K(W)\varepsilon}{\|m_j - w_i\|^2}.
\end{aligned}$$

This allows us conclude that

$$\|\widetilde{m}_{ij} - w_i\| = \delta_{ij}\|m_j - w_i\| \leq \frac{8K(W)\varepsilon}{\|m_j - w_i\|} \leq \frac{32K(W)\varepsilon}{\sigma_k(W)}, \tag{3.3}$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

because if $m_j = w_j$, then $\|m_j - w_i\| \geq \sqrt{2}\sigma_k(W) \geq \sigma_k(W)/4$, and if $m_j = q_s$, then

$$\|m_j - w_i\| \geq \|w_i\| - \|m_j\| \geq K(W) - 2\varepsilon - K(W)/2 \geq K(W)/2 - \sigma_k(W)\mathcal{K}(W)^{-1}/4 \geq \sigma_k(W)/4.$$

Finally, by (3.1),

$$\min_i \|x - w_i\| \leq \varepsilon + \max_{i:\mathscr{S}_i \neq \emptyset} \max_{j \neq i} \|\widetilde{m}_{ij} - w_i\| \leq \varepsilon + \frac{32K(W)\varepsilon}{\sigma_k(W)} \leq 33\mathcal{K}(W)\varepsilon. \qquad \square$$

The consequences of Theorem 3 are the following.

- The first step of SPA is more robust than the next ones. The reason is that the projection step might lead to an increase of the error bound by a factor of $\mathcal{K}(W)$, so that Theorem 1 becomes tight after the first step of SPA. We will prove this in section 5 with an example. From a practical point of view, this is interesting information: we know that the first extracted vertex by SPA is more reliable than the next ones.
- When $H^\top e = e$, one can use T-SPA. Since T-SPA only translates at the first step while increasing the conditioning by a factor at most two (see section 2.1), the first two steps of T-SPA follow the bound of Theorem 3 and are therefore more robust than the next ones.
- In the proof of Theorem 3, the assumption that H is separable can be relaxed: we only need the columns of W to appear as columns of X (up to the noise), not that of Q. In other words, only $H(1:k,:)$ needs to be separable. We will use this observation in the proof of the Theorem 4 in the next section.
- Since the only hypothesis on Q is a bound on the norm of its columns, Theorem 3 generalizes easily to substochastic matrices H (that is, $H^\top e \leq e$). In fact, we only need to add a zero column to Q and a row to the substochastic H to extend it to a stochastic matrix.

3.2. SPA when $r = 2$. Another less direct consequence of Theorem 3 is that we can solve rank-2 separable NMF with better guarantees, with a bound on the noise of $\varepsilon \leq \mathcal{O}(\frac{\sigma_r(W)}{\mathcal{K}(W)})$, and an error bound in $\mathcal{O}(\varepsilon\mathcal{K}(W))$.

Unfortunately, we show in the next section that the following steps of SPA when $r \geq 3$ cannot be as good and that the bound of Theorem 1 is asymptotically tight.

THEOREM 4. *Let $W \in \mathbb{R}^{m\times 2}$, $Q \in \mathbb{R}^{m\times(r-2)}$ and $X = [W\ Q]H + N \in \mathbb{R}^{m\times n}$ with $2 \leq m$, where W is full rank, $K(N) < \varepsilon \leq \mathcal{O}(\sigma_2(W)\mathcal{K}(W)^{-1})$, $K(Q) \leq \sigma_2(W)/2$, and $H \in \mathbb{R}_+^{r\times n}$ has stochastic columns and two of its columns are e_1, e_2. Suppose x_1, x_2 are the two columns chosen in order by the SPA algorithm applied on X. Then we have that*

$$\min_\pi \max_{i=1,2} \|x_i - w_{\pi(i)}\| \leq \mathcal{O}(\mathcal{K}(W)\varepsilon),$$

where π is either the identity function or the transposition $\pi(i) = 3 - i$.

Proof. Suppose $\varepsilon \leq \sigma_2(W)\mathcal{K}(W)^{-1}/305$. Since $1/305 \leq 1/8$, we can apply Theorem 3 and find that, up to a permutation of the columns of W and H, SPA chooses at the first step the column x_1 and $\|r_1\| \leq 33\mathcal{K}(W)\varepsilon$, where $r_1 := x_1 - w_1$. Notice that since x_1 is the column of X with the largest norm and H has among its columns the vectors e_1, e_2, then $w_i + n_i$ are columns of X and in particular

$$\|x_1\| \geq \max_i \|w_i + n_i\| \geq \max_i \|w_i\| - \varepsilon = K(W) - \varepsilon.$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Let $P_v := I - vv^\top/\|v\|^2$ be the orthogonal projection to the hyperplane orthogonal to the vector v, and let $c = 1/305$.

$$\begin{aligned}
\|P_{x_1} - P_{w_1}\| &= \left\| \frac{x_1 x_1^\top}{\|x_1\|^2} - \frac{w_1 w_1^\top}{\|w_1\|^2} \right\| = \left\| \frac{(r_1+w_1)(r_1+w_1)^\top}{\|x_1\|^2} - \frac{w_1 w_1^\top}{\|w_1\|^2} \right\| \\
&= \left\| \frac{r_1 r_1^\top + r_1 w_1^\top + w_1 r_1^\top}{\|x_1\|^2} + w_1 w_1^\top \left(\frac{1}{\|x_1\|^2} - \frac{1}{\|w_1\|^2} \right) \right\| \\
&\le \frac{\|r_1\|^2 + 2\|r_1\|\|w_1\|}{\|x_1\|^2} + \frac{|\|w_1\|^2 - \|x_1\|^2|}{\|x_1\|^2} \\
&\le \frac{\|r_1\|^2 + 2\|r_1\|\|w_1\|}{\|x_1\|^2} + \frac{2\|w_1\|\|r_1\| + \|r_1\|^2}{\|x_1\|^2} = 2\|r\| \frac{\|r_1\| + 2\|w_1\|}{\|x_1\|^2} \\
&\le 2\|r_1\| \frac{33\mathcal{K}(W)\varepsilon + 2K(W)}{(K(W)-\varepsilon)^2} \le 2\|r_1\| \frac{33c\sigma_2(W) + 2K(W)}{(K(W) - c\sigma_2(W))^2} \\
&\le \frac{2\|r_1\|}{K(W)} \frac{33c+2}{(1-c)^2} \le \frac{\mathcal{K}(W)\varepsilon}{K(W)} \gamma(c),
\end{aligned}$$

where $\gamma(c) = 66\frac{33c+2}{(1-c)^2} \le 141$. As a consequence,

$$\begin{aligned}
\|P_{x_1} w_2\| &\ge \|P_{w_1} w_2\| - \|(P_{x_1} - P_{w_1}) w_2\| \ge \sigma_2(W) - \|P_{x_1} - P_{w_1}\|\|w_2\| \\
&\ge \sigma_2(W) - \frac{\mathcal{K}(W)\varepsilon}{K(W)} \gamma(c) \|w_2\| \ge (1 - c\gamma(c))\sigma_2(W) \ge \frac{1 - c\gamma(c)}{c} \mathcal{K}(W)\varepsilon.
\end{aligned}$$

If we now let $\widetilde{X} := P_{x_1} X = P_{x_1}[W\ Q]H + P_{x_1} N =: [\widetilde{W}\ \widetilde{Q}]H + \widetilde{N}$, then the second step SPA identifies $\widetilde{x}_2$, the column of $\widetilde{X}$ with the largest norm. Notice that $\widetilde{w}_2 = P_{x_1} w_2$ and

$$\begin{aligned}
\|\widetilde{w}_1\| = \|P_{x_1} w_1\| &\le \|x_1 - w_1\| \le 33\mathcal{K}(W)\varepsilon \le \frac{33c}{1 - c\gamma(c)} \|\widetilde{w}_2\| \\
&\le \frac{1}{2} \frac{1}{1 - c\gamma(c)} \|\widetilde{w}_2\| < \|\widetilde{w}_2\|.
\end{aligned}$$

Notice moreover that $K(\widetilde{N}) \le K(N) \le \varepsilon$ and

$$K(\widetilde{Q}) \le K(Q) \le \frac{\sigma_2(W)}{2} \le \frac{1}{2} \frac{1}{1 - c\gamma(c)} \|\widetilde{w}_2\| < \|\widetilde{w}_2\|.$$

We can write $\widetilde{x}_2 = \lambda \widetilde{w} + (1-\lambda)\widetilde{w}_2 + \widetilde{n}_2$ for a $0 \le \lambda \le 1$, for some $\widetilde{w}$ in the convex hull of $[\widetilde{Q}\ \widetilde{w}_1]$. Since

$$\|\widetilde{w}\| \le \frac{1}{2} \frac{1}{1 - c\gamma(c)} \|\widetilde{w}_2\| < \|\widetilde{w}_2\|$$

and since H has e_2 as a column, then one of the columns of $\widetilde{X}$ is equal to $\widetilde{w}_2 + \widetilde{n}$ for some column $\widetilde{n}$ of $\widetilde{N}$, but $\widetilde{x}_2$ is its column with largest norm, so

$$\begin{aligned}
&\|\widetilde{w}_2\| - \varepsilon \le \|\widetilde{w}_2 + \widetilde{n}\| \le \|\widetilde{x}_2\| \le \lambda\|\widetilde{w}\| + (1-\lambda)\|\widetilde{w}_2\| + \varepsilon \implies \lambda(\|\widetilde{w}_2\| - \|\widetilde{w}\|) \le 2\varepsilon \\
&\implies \lambda \le \frac{2\varepsilon}{\|\widetilde{w}_2\| - \|\widetilde{w}\|} \le \frac{2\varepsilon}{(1 - \frac{1}{2}\frac{1}{1-c\gamma(c)})\|\widetilde{w}_2\|}.
\end{aligned}$$

Since the projection is linear, we have $x_2 = \lambda w + (1-\lambda) w_2 + n_2$ for some w in the convex hull of $[Q\ w_1]$ and

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$$\begin{aligned}\|x_2 - w_2\| &\leq \varepsilon + \lambda\|w - w_2\| \leq \varepsilon + \frac{2\varepsilon\|w_1 - w_2\|}{(1 - \frac{1}{2}\frac{1}{1-c\gamma(c)})\|\widetilde{w}_2\|} \\ &\leq \varepsilon + \frac{4K(W)}{(1 - \frac{1}{2}\frac{1}{1-c\gamma(c)})(1 - c\gamma(c))\sigma_2(W)}\varepsilon \\ &= (1 + \beta(c)\mathcal{K}(W))\,\varepsilon \leq (1 + \beta(c))\,\mathcal{K}(W)\varepsilon,\end{aligned}$$

where

$$\beta(c) = \frac{4}{(1 - \frac{1}{2}\frac{1}{1-c\gamma(c)})(1 - c\gamma(c))} = \frac{4}{\frac{1}{2} - c\gamma(c)} \leq 99. \qquad \square$$

Remark 3. Theorem 4 does not give an explicit relation between the constants in the noise bound $\varepsilon \leq \mathcal{O}(\sigma_2(W)\mathcal{K}(W)^{-1})$ and the constants of the error bound $\min_{\pi\in S^2}\max_{i=1,2}\|x_i - w_{\pi(i)}\| \leq \mathcal{O}(\mathcal{K}(W)\varepsilon)$ as in Theorem 3. Carefully retracing the steps in the proof, we can find that if $\varepsilon \leq c\sigma_2(W)\mathcal{K}(W)^{-1}$ with $c \leq 1/150$, then the error in the first and second step of SPA are, respectively, bounded by

$$\left[1 + \frac{16}{1-4c}\right]\mathcal{K}(W)\varepsilon \leq 18\mathcal{K}(W)\varepsilon, \qquad \left[1 + \frac{8(1-4c)^2(1-c)^2}{1 - 146c - 547c^2 + 376c^3 - 48c^4}\right]\mathcal{K}(W)\varepsilon.$$

If, for example, we want to keep the same error $33\mathcal{K}(W)\varepsilon$ as in Theorem 3, we need $c \leq 1/196$. If instead we want to match the error bound $18\mathcal{K}(W)\varepsilon$ of the first step, we would need $c \leq 1/271$.
Note that the second formula above, for the second step error, can be simplified using the inequality

$$\left[1 + \frac{8(1-4c)^2(1-c)^2}{1 - 146c - 547c^2 + 376c^3 - 48c^4}\right] \leq 9 + 4400c$$

that holds for any $c \leq 1/200$.

4. Improved bounds for the first two steps of T-SPA and when $r = 3$.

T-SPA is the same algorithm as SPA, except at the first step when it performs a translation instead of an orthogonal projection. Because of that, the bounds for T-SPA will be "one step ahead" of that of SPA. More precisely, we show improved bounds for the first two steps of T-SPA in section 4.1 (Theorem 5) and when $r = 3$ in section 4.2 (Theorem 6).

4.1. First two steps of T-SPA.

Similarly as done for the first step of SPA in section 3.1, we prove in this section that T-SPA has bounds on the noise of $\mathcal{O}(\sigma_r(W)/\mathcal{K}(W))$ and error bounds of $\mathcal{O}(\varepsilon\mathcal{K}(W))$ for the first two steps.

THEOREM 5. *Let $W \in \mathbb{R}^{m\times r}$ and $X = WH + N \in \mathbb{R}^{m\times n}$ with $r \leq \min\{m, n\}$, where W is full rank, $K(N) < \varepsilon \leq \mathcal{O}(\sigma_r(W)\mathcal{K}(W)^{-1})$, and $H \in \mathbb{R}_+^{r\times n}$ is separable and has stochastic columns. Given x_1, x_2 the first two columns chosen by the T-SPA algorithm applied on X, we have that*

$$\min_{\pi\in S^r}\max_{i=1,2}\|x_i - w_{\pi(i)}\| \leq \mathcal{O}(\mathcal{K}(W)\varepsilon),$$

where S^r is the set of permutations of $\{1, 2, \ldots, r\}$.

Before proving Theorem 5, we need two auxiliary lemmas.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

LEMMA 3. *Let $W \in \mathbb{R}^{m\times r}$, where $r \le m$ and W is full rank. Let $v \in \operatorname{conv}(W)$ be such that in the convex combination $v = \sum_i \lambda_i w_i$ we have $\lambda_r \ge 1/2$. If $\hat{W}$ are the first $r-1$ columns of the matrix $W - ve^\top$, then*

$$\sigma_r(W) \le \frac{2}{2-\sqrt{2}}\sigma_{r-1}(\hat{W}), \qquad K(\hat{W}) \le 2K(W).$$

Proof. See Appendix A.3. ∎

LEMMA 4. *Let $W \in \mathbb{R}^{m\times r}$ and $X = WH + N \in \mathbb{R}^{m\times n}$ with $r \le \min\{m,n\}$, where W is full rank, $K(N) < \varepsilon \le c\sigma_r(W)\mathcal{K}(W)^{-1}$, $c \le 1/233$, and $H \in \mathbb{R}_+^{r\times n}$ is separable and has stochastic columns. Suppose x_1 is the first column chosen by the T-SPA when applied on X, and it is associated to w_1 in the sense of Theorem 3, i.e., $x_1 = Wh_1 + n_1$, where $h_{1,1} \ge 1/2$. Call $w := Wh_1$ and let*

$$\hat{X} := X - x_1 e^\top = (W - we^\top)H + (N - n_1 e^\top) := [r_1\ \hat{W}]H + \widetilde{N},$$

where $r_1 := w_1 - w$. Then $\hat{W}$ is full rank and

- $\mathcal{K}(\hat{W}) \le \frac{4}{2-\sqrt{2}}\mathcal{K}(W)$,
- $\sigma_r(W)\mathcal{K}(W)^{-1} \le \frac{4}{3-2\sqrt{2}}\sigma_{r-1}(\hat{W})\mathcal{K}(\hat{W})^{-1}$,
- $\|r_1\| \le \frac{1}{2}\sigma_{r-1}(\hat{W})$.

Proof. See Appendix A.4. ∎

We can now prove Theorem 5.

Proof of Theorem 5. Up to a permutation of the columns of X, W, and H, T-SPA chooses at the first step the column x_1 and it is associated to w_1 in the sense of Theorem 3, i.e., $x_1 = Wh_1 + n_1$, where $h_{1,1} \ge 1/2$. Call $w := Wh_1$ and let

$$\hat{X} := X - x_1 e^\top = (W - we^\top)H + (N - n_1 e^\top) := [r_1\ \hat{W}]H + \widetilde{N},$$

where $r_1 := w_1 - w$ and $K(\widetilde{N}) < 2\varepsilon$. Suppose $\varepsilon \le c\sigma_r(W)\mathcal{K}(W)^{-1}$ with $c = 1/374$. Lemma 4 tells us that $\hat{W}$ is full rank, $\|r_1\| \le \frac{1}{2}\sigma_{r-1}(\hat{W}) \le \frac{1}{2}K(\hat{W})$, and

$$\varepsilon \le c\sigma_r(W)\mathcal{K}(W)^{-1} \le \frac{4c}{3-2\sqrt{2}}\sigma_{r-1}(\hat{W})\mathcal{K}(\hat{W})^{-1} \le \sigma_{r-1}(\hat{W})\mathcal{K}(\hat{W})^{-1}/16.$$

We can thus use again Theorem 3 to obtain that, up to permutations, the second step of T-SPA extracts the column $\hat{x}_2 = x_2 - x_1$ of $\hat{X}$ and it is close to the first column of $\hat{W}$, that is, $\hat{w}_1 = w_2 - w$, in the sense $\|\hat{x}_2 - \hat{w}_1\| \le 66\mathcal{K}(\hat{W})\varepsilon$. This lets us conclude that x_2 is close to w_2 as, by Lemma 4,

$$\|x_2 - w_2\| \le \|\hat{x}_2 - \hat{w}_1\| + \|x_1 - w\| \le 67\mathcal{K}(\hat{W})\varepsilon \le 67\frac{4}{2-\sqrt{2}}\mathcal{K}(W)\varepsilon. \qquad ∎$$

4.2. T-SPA when $r = 3$. As done for SPA when $r = 2$ in section 3.2, we prove in this section that when $r = 3$, T-SPA has a bound on the noise of $\varepsilon \le \mathcal{O}(\frac{\sigma_r(W)}{\mathcal{K}(W)})$ and an error bound in $\mathcal{O}(\varepsilon\mathcal{K}(W))$.

THEOREM 6. *Let $W \in \mathbb{R}^{m\times 3}$ and $X = WH + N \in \mathbb{R}^{m\times n}$ with $3 \le \min\{m,n\}$, where W is full rank, $K(N) < \varepsilon \le \mathcal{O}(\sigma_3(W)\mathcal{K}(W)^{-1})$, and $H \in \mathbb{R}_+^{3\times n}$ is separable and has stochastic columns. Given x_1, x_2, x_3 the three columns chosen by the T-SPA algorithm applied on X, we have that*

$$\min_{\pi\in S^3}\max_{i=1,2,3}\|x_i - w_{\pi(i)}\| \le \mathcal{O}(\mathcal{K}(W)\varepsilon).$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Proof. Up to a permutation of the columns of X, W, and H, T-SPA chooses at the first step the column x_1 and it is associated to w_1 in the sense of Theorem 3, i.e., $x_1 = Wh_1 + n_1$, where $h_{1,1} \geq 1/2$. Call $w := Wh_1$ and let

$$\hat{X} := X - x_1 e^\top = (W - we^\top)H + (N - n_1 e^\top) := [r_1 \ \hat{W}]H + \widetilde{N},$$

where $r_1 := w_1 - w$ and $K(\widetilde{N}) < 2\varepsilon$. Suppose $\varepsilon \leq c\sigma_3(W)\mathcal{K}(W)^{-1}$ with $c \leq 1/233$, so that we can apply Lemma 4, which tells us that $\hat{W}$ is full rank, $\|r_1\| \leq \frac{1}{2}\sigma_2(\hat{W})$, and

$$\varepsilon \leq c\sigma_3(W)\mathcal{K}(W)^{-1} \leq \frac{4c}{3-2\sqrt{2}}\sigma_2(\hat{W})\mathcal{K}(\hat{W})^{-1} \leq \mathcal{O}(\sigma_2(\hat{W})\mathcal{K}(\hat{W})^{-1}).$$

We can thus use Theorem 4 to obtain that, up to permutations, the remaining steps of T-SPA extract the columns $\hat{x}_2, \hat{x}_3$ of $\hat{X}$ and they are close to the columns of $\hat{W}$, which are $\hat{w}_1 = w_2 - w$ and $\hat{w}_2 = w_3 - w$, in the sense $\|\hat{x}_2 - \hat{w}_1\| \leq \mathcal{O}(\mathcal{K}(\hat{W})\varepsilon)$ and $\|\hat{x}_3 - \hat{w}_2\| \leq \mathcal{O}(\mathcal{K}(\hat{W})\varepsilon)$. We thus conclude that

$$\max_{i=2,3} \|x_i - w_i\| \leq \max_{i=2,3} \|\hat{x}_i - \hat{w}_{i-1}\| + \|x_1 - w\| \leq \mathcal{O}(\mathcal{K}(\hat{W})\varepsilon) + \varepsilon = \mathcal{O}(\mathcal{K}(\hat{W})\varepsilon). \quad \square$$

In a sense, Theorems 4 and 6 say the same thing: Theorem 4 shows that if we have some points that are approximately convex combinations of two vertices and the origin, then we are able to extract two vertices; Theorem 6 says that if we have points that are approximately convex combinations of three vertices, then once we bring the first vertex to zero, we can extract the other two.

5. Tightness of the error bound in Theorem 1 when $r \geq 3$. In this section, we prove that the error bound in Theorem 1 is tight when $r \geq 3$ (up to the hidden constant factors), which was an important open question regarding SPA [15, p. 231].

THEOREM 7. *For $r \geq 3$, the second step of SPA might have an error in $\Omega(\varepsilon\mathcal{K}^2(W))$. In particular, there exist a full rank matrix $W \in \mathbb{R}^{3\times 3}$ and an absolute constant $c > 0$ such that, for any ε in the interval $(0, c)$, there exist $H_\varepsilon \in \mathbb{R}^{3\times 4}$ separable and column stochastic, and $N_\varepsilon \in \mathbb{R}^{3\times 4}$ with $K(N_\varepsilon) \leq \varepsilon$, for which the error of the second step of SPA applied to $WH_\varepsilon + N_\varepsilon$ is larger than $c\varepsilon\mathcal{K}^2(W)$. This implies that the error bound in Theorem 1 is tight for $r \geq 3$, up to constant terms.*

Proof. Let us provide a family of examples where SPA achieves the worst-case bound, already at the second step.

Let $0 < \delta < 1/4$, $K > 1$, $0 < \varepsilon < K\delta^2/3$, and

$$W = K\begin{pmatrix} 1 & 0 & 1/2 \\ 0 & 1/2 & 1/2 \\ 0 & -\delta & \delta \end{pmatrix}.$$

The product of the singular values of W is $|\det(W)| = K^3\delta$ and $K(W) = K \leq \|W\|$ since $1/4 + \delta^2 < 1/2 + \delta^2 < 1$. Given the decomposition

$$W = W_1 + W_2 := K\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1/2 & 0 \\ 0 & -\delta & 0 \end{pmatrix} + K\begin{pmatrix} 0 & 0 & 1/2 \\ 0 & 0 & 1/2 \\ 0 & 0 & \delta \end{pmatrix},$$

it is immediate to see that $\sigma_1(W_1) = K$, $\sigma_2(W_1) = K\sqrt{\delta^2 + 1/4} > K/2$, and $\|W_2\| = K\sqrt{\delta^2 + 1/2} < K$, so by triangular inequality $\sigma_2(W) \leq \sigma_1(W) = \|W\| \leq \|W_1\| +$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$\|W_2\| < 2K$ and by the interlacing property of the singular values, $\sigma_1(W) \geq \sigma_2(W) \geq \sigma_2(W_1) > K/2$. As a consequence,

$$\sigma_3(W) = \frac{K^3\delta}{\sigma_1(W)\sigma_2(W)} \in \left[\frac{K\delta}{4}, 4K\delta\right] \implies \frac{8}{\delta} \geq \mathcal{K}(W) \geq \frac{1}{8\delta}.$$

Notice that

$$\frac{1}{2} + 2\frac{\varepsilon}{K} < \frac{1}{2} + \frac{2}{3}\delta^2 = \sqrt{\frac{1}{4} + \frac{2}{3}\delta^2 + \frac{4}{9}\delta^4} < \sqrt{\frac{1}{4} + \frac{2}{3}\delta^2 + \frac{1}{36}\delta^2} < \sqrt{\frac{1}{4} + \delta^2}$$

so there exists γ such that $\delta > \gamma > 0$ and $\sqrt{1/4 + \delta^2} - 2\varepsilon/K = \sqrt{1/4 + \gamma^2}$. Let also

$$\widetilde{v} := K\begin{pmatrix} 0 \\ 1/2 \\ -\delta \end{pmatrix} + \frac{1}{2}\left(1 - \frac{\gamma}{\delta}\right) K \begin{pmatrix} 1/2 \\ 0 \\ 2\delta \end{pmatrix} = K\begin{pmatrix} 1/4\left(1 - \frac{\gamma}{\delta}\right) \\ 1/2 \\ -\gamma \end{pmatrix},$$

$$v := \widetilde{v} + \frac{\varepsilon}{\sqrt{\frac{1}{4} + \gamma^2}}\begin{pmatrix} 0 \\ 1/2 \\ -\gamma \end{pmatrix}.$$

The vector $\widetilde{v}$ is in the convex hull of W since $\widetilde{v} = w_3(1-\gamma/\delta)/2 + w_2[1-(1-\gamma/\delta)/2]$, and v is a ε-perturbation of it. Notice moreover that

$$\|v\| \leq \varepsilon + \|\widetilde{v}\| \leq \varepsilon + K\sqrt{\frac{1}{16} + \frac{1}{4} + \gamma^2} < K\frac{\delta^2}{3} + K\sqrt{\frac{5}{16} + \delta^2} < K\left(\frac{1}{3 \cdot 16} + \sqrt{\frac{6}{16}}\right) < K.$$

Let

$$z_1 := w_2 - \frac{\varepsilon}{\sqrt{\frac{1}{4} + \delta^2}}\begin{pmatrix} 0 \\ 1/2 \\ -\delta \end{pmatrix} = \left(-\frac{\varepsilon}{\sqrt{\frac{1}{4} + \delta^2}} + K\right)\begin{pmatrix} 0 \\ 1/2 \\ -\delta \end{pmatrix},$$

which is the second column of W up to a perturbation of norm ε. In particular, $\|z_1\| = \|w_2\| - \varepsilon < K - \varepsilon$. Finally, let

$$z_2 := w_3 - \frac{\varepsilon}{\sqrt{\frac{1}{4} + \delta^2}}\begin{pmatrix} 0 \\ 1/2 \\ \delta \end{pmatrix} = \left(-\frac{\varepsilon}{\sqrt{\frac{1}{4} + \delta^2}} + K\right)\begin{pmatrix} 0 \\ 1/2 \\ \delta \end{pmatrix} + K\begin{pmatrix} 1/2 \\ 0 \\ 0 \end{pmatrix},$$

which is the third column of W up to a perturbation of norm ε. In particular,

$$\begin{aligned}\|z_2\| \leq \|w_3\| + \varepsilon &= K\sqrt{\frac{1}{2} + \delta^2} + 2\varepsilon - \varepsilon < K\sqrt{\frac{1}{2} + \frac{1}{16}} + \frac{2}{3}K\delta^2 - \varepsilon \\ &< K\left(\frac{3}{4} + \frac{1}{24}\right) - \varepsilon < K - \varepsilon.\end{aligned}$$

Let us build the matrix

$$X = \begin{pmatrix} K & & & \\ 0 & z_1 & z_2 & v \\ 0 & & & \end{pmatrix},$$

which is in the form $WH + N$ with H separable and column stochastic, W full rank, and $K(N) \leq \varepsilon$.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

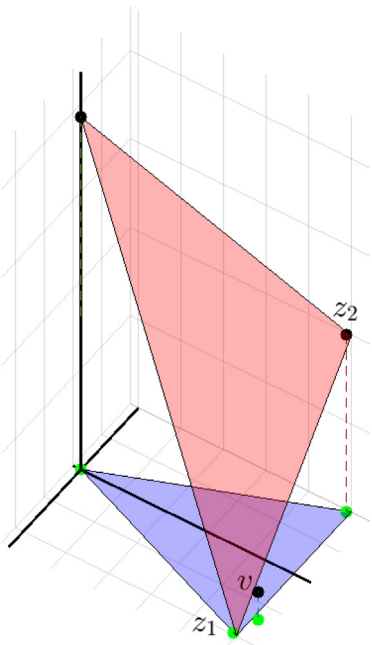

FIG. 1. *Columns of matrix X as black dots and of matrix $\widetilde{X}$ as green dots. (Color images are available online.)*

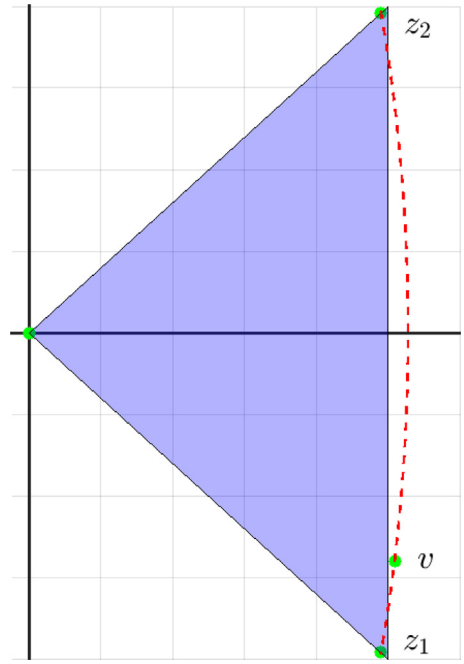

FIG. 2. *Columns of matrix $\widetilde{X}(\{2,3\},:)$. The points of the dashed red line have the same norm. (Color images are available online.)*

The first step of SPA identifies the first column as the vector of largest norm and projects on its orthogonal space, obtaining the new residual matrix

$$\widetilde{X} = K \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1/2\left(1-\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\delta^2}}\right) & 1/2\left(1-\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\delta^2}}\right) & 1/2\left(1+\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\gamma^2}}\right) \\ 0 & -\delta\left(1-\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\delta^2}}\right) & \delta\left(1-\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\delta^2}}\right) & -\gamma\left(1+\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\gamma^2}}\right) \end{pmatrix}.$$

Figure 1 represents the columns of X and $\widetilde{X}$, and Figure 2 that of $\widetilde{X}$ discarding the first coordinate which is equal to zero for all columns.

Notice that all the columns of $\widetilde{X}$ have the same norm since

$$\left(\frac{1}{4}+\gamma^2\right)\left(\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\gamma^2}}+1\right)^2 = \left(\frac{\varepsilon}{K}+\sqrt{\frac{1}{4}+\gamma^2}\right)^2$$
$$= \left(-\frac{\varepsilon}{K}+\sqrt{\frac{1}{4}+\delta^2}\right)^2 = \left(\frac{1}{4}+\delta^2\right)\left(-\frac{\varepsilon}{K\sqrt{\frac{1}{4}+\delta^2}}+1\right)^2.$$

Hence SPA can select the last one, v, at the second step. In that case, the final error of SPA will be larger than $\min_i \|w_i - v\|$. Moreover, recall that v is $\widetilde{v}$ plus a perturbation of norm ε, so $\|w_i - v\| \geq \|w_i - \widetilde{v}\| - \varepsilon$. First, $\|w_1 - \widetilde{v}\| \geq 3K/4 > 9\varepsilon/(4\delta^2)$ and $\|w_3 - \widetilde{v}\| \geq K/4 > 3\varepsilon/(4\delta^2)$. For the second one, we get

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$$\|w_2 - \widetilde{v}\| \geq K\frac{\delta-\gamma}{4\delta} = K\frac{\delta^2-\gamma^2}{4\delta(\delta+\gamma)} \geq K\frac{\left(\frac{1}{4}+\delta^2\right)-\left(\frac{1}{4}+\gamma^2\right)}{8\delta^2}$$
$$= K\frac{\sqrt{\frac{1}{4}+\delta^2}-\sqrt{\frac{1}{4}+\gamma^2}}{8\delta^2}\left(\sqrt{\frac{1}{4}+\delta^2}+\sqrt{\frac{1}{4}+\gamma^2}\right) \geq K\frac{2\varepsilon/K}{8\delta^2} = \frac{\varepsilon}{4\delta^2}.$$

As a consequence, the error will be larger than

$$\frac{\varepsilon}{4\delta^2} - \varepsilon \geq \frac{3\varepsilon}{16\delta^2} \geq \frac{3\varepsilon}{16}\frac{\mathcal{K}^2(W)}{64}. \qquad \square$$

6. Tightness of preconditioned SPA variants. In this section, we show that the two preconditioned variants of SPA described in section 2.2, namely, SPA^2 and MVE-SPA, have tight error bounds. We do this again with examples achieving the worst-case bounds.

6.1. Tightness of SPA^2. Recall the SPA^2 first runs SPA on X to obtain $W = X(:,\mathcal{J})$ and then reruns SPA on $W^\dagger X$ to obtain $W' = X(:,\mathcal{J}')$. It turns out this strategy reduces the error bound of SPA to $\mathcal{O}(\varepsilon\mathcal{K}(W))$.

THEOREM 8 ([18, Theorem 4.1]). *Under the same assumptions and bounds on the noise as Theorem* 1*, the index set $\mathcal{J}'$ extracted by SPA^2 satisfies the following error bound:*

$$\min_{\pi\in S^r}\max_{1\leq k\leq r}\|W(:,k) - X(:,\mathcal{J}'_{\pi(k)})\| \leq \mathcal{O}\left(\varepsilon\mathcal{K}(W)\right).$$

Let us prove that this bound is tight, as soon as $r \geq 2$.

THEOREM 9. *The error bound of Theorem* 8 *is tight for $r \geq 2$.*

Proof. Let M be any positive real number and let $0 < \delta < 1/2$, $0 < \varepsilon < \delta M/6$, $\alpha = \varepsilon^2/(6\delta M) < \varepsilon/36$. Consider

$$W = M\begin{pmatrix}1 & 0\\ 0 & \delta\end{pmatrix},$$
$$X = \begin{pmatrix}M & 0 & \frac{2\varepsilon+\alpha}{\delta}\\ 0 & \delta M-\varepsilon & \delta M-\varepsilon-\alpha\end{pmatrix} = W\begin{pmatrix}1 & 0 & z\\ 0 & 1 & 1-z\end{pmatrix} + \begin{pmatrix}0 & 0 & 0\\ 0 & -\varepsilon & \varepsilon\end{pmatrix},$$

where

$$0 < z = \frac{2\varepsilon+\alpha}{\delta M} < \frac{3\varepsilon}{\delta M} < \frac{1}{2}.$$

The columns of X and W are represented in Figure 3. Since $0 < \delta M - \varepsilon < M$, the squared norm of the last column of X is

$$\left(\frac{2\varepsilon+\alpha}{\delta}\right)^2 + (\delta M-\varepsilon-\alpha)^2 < \left(\frac{3\varepsilon}{\delta}\right)^2 + (\delta M)^2 < \left(\frac{M}{2}\right)^2 + \left(\frac{M}{2}\right)^2 < M^2.$$

SPA thus chooses first the first index, and then the second one since $\delta M - \varepsilon > \delta M - \varepsilon - \alpha$. We get the preconditioner

$$\widetilde{W} = \begin{pmatrix}M & 0\\ 0 & \delta M-\varepsilon\end{pmatrix} \Longrightarrow \widetilde{X} = \widetilde{W}^{-1}X = \begin{pmatrix}1 & 0 & \frac{2\varepsilon+\alpha}{\delta M}\\ 0 & 1 & 1-\frac{\alpha}{\delta M-\varepsilon}\end{pmatrix}.$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

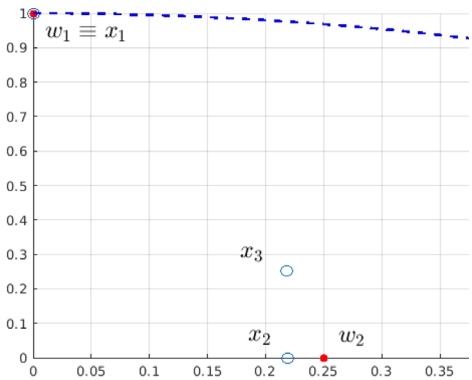

FIG. 3. *Columns of matrix X as empty blue circles and columns of matrix W as red dots. The points on the circle (dashed blue line) have the same norm.*

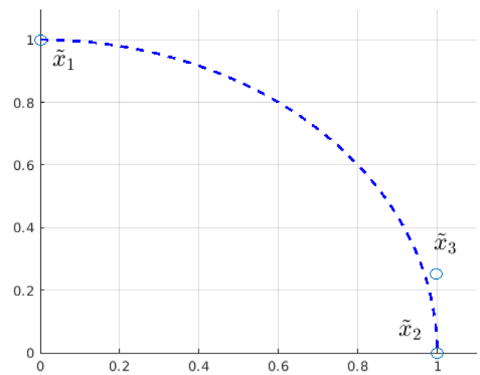

FIG. 4. *Columns of the preconditioned matrix $\widetilde{X}$. The points on the circle (dashed blue line) have the same norm.*

The columns of $\widetilde{X}$ are represented in Figure 4. Since $\delta M - \varepsilon > 5\delta M/6$, then the norm squared of the last column of $\widetilde{X}$ is now

$$\left(\frac{2\varepsilon+\alpha}{\delta M}\right)^2 + \left(1 - \frac{\alpha}{\delta M - \varepsilon}\right)^2 > \frac{4\varepsilon^2}{(\delta M)^2} + 1 - \frac{2\alpha}{\delta M - \varepsilon} > \frac{4\varepsilon^2}{(\delta M)^2} + 1 - \frac{1}{5}\frac{2\varepsilon^2}{(\delta M)^2} > 1,$$

so the first index chosen by the algorithm SPA^2 is the third one. This means that the final error of SPA^2 will be at least the minimum distance between the last column of X and the columns of W:

$$\begin{aligned}(\text{error of SPA}^2)^2 &\geq \min\left\{\left(\frac{2\varepsilon+\alpha}{\delta} - M\right)^2 + (\delta M - \varepsilon - \alpha)^2, \left(\frac{2\varepsilon+\alpha}{\delta}\right)^2 + \alpha^2\right\} \\ &\geq \min\left\{\left(M - \frac{2\varepsilon+\alpha}{\delta}\right)^2, \left(\frac{2\varepsilon}{\delta}\right)^2\right\} \geq \min\left\{\left(M - \frac{3\varepsilon}{\delta}\right)^2, \left(\frac{2\varepsilon}{\delta}\right)^2\right\} \\ &\geq \min\left\{\left(\frac{3\varepsilon}{\delta}\right)^2, \left(\frac{2\varepsilon}{\delta}\right)^2\right\} = \left(\frac{2\varepsilon}{\delta}\right)^2.\end{aligned}$$

Since $\mathcal{K}(W) = 1/\delta$, we obtain that the error bound of SPA^2 must be larger than $2\mathcal{K}(W)\varepsilon$. ∎

6.2. Tightness of MVE-SPA. MVE-SPA needs to be applied on an r-by-n matrix, so the input matrix needs to be replaced by a low-rank approximation (e.g., using the truncated SVD) as explained in section 2.2. Then MVE-SPA first identifies the MVE centered at the origin and containing all data points, defined as $\{v \in \mathbb{R}^r \mid v^\top A v \leq 1\}$, where $A \succ 0$, and then applies SPA on $\widetilde{X} = A^{1/2}X$ to obtain an index set $\mathcal{J}$. After the preconditioning, all data points previously on the border of the ellipsoid have exactly norm 1, whereas all the other points have norm strictly less than 1. As a consequence, it can be shown that MVE-SPA can extract any index corresponding to data points on the border of the ellipsoid, that is, $\{v \in \mathbb{R}^r \mid v^\top A v = 1\}$ [20]; see also [28].

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

THEOREM 10 ([20, Theorem 2.9]). *Let $X = WH + N \in \mathbb{R}^{r\times n}$, where* $\mathrm{rk}(W) = r$, *$H \in \mathbb{R}_+^{r\times n}$ is a separable matrix satisfying $H^\top e \le e$, and*

$$\varepsilon = K(N) \le \mathcal{O}\left(\frac{\sigma_r(W)}{r\sqrt{r}}\right).$$

Then MVE-SPA returns a set of indices $\mathcal{J}$ such that

$$\min_{\pi\in S^r} \max_{1\le k\le r} \|W(:,k) - X(:,\mathcal{J}_{\pi(k)})\| \le \mathcal{O}\left(\varepsilon\mathcal{K}(W)\right).$$

THEOREM 11. *The error bound in Theorem* 10 *is tight.*

Proof. We can use a similar example as for SPA^2 in Theorem 8. Let M be any positive real number and let $0 < \delta < 1$, $0 < \varepsilon < \delta M/4$. Consider

$$W = M\begin{pmatrix} 1 & 0 \\ 0 & \delta \end{pmatrix},$$
$$X = \begin{pmatrix} M & 0 & \frac{2\varepsilon}{\delta} \\ 0 & \delta M - \varepsilon & \delta M - \varepsilon \end{pmatrix} = W\begin{pmatrix} 1 & 0 & z \\ 0 & 1 & 1-z \end{pmatrix} + \begin{pmatrix} 0 & 0 & 0 \\ 0 & -\varepsilon & \varepsilon \end{pmatrix},$$

where $z = 2\varepsilon/(M\delta) < 1/2$. The minimum volume ellipsoid going through the first two columns of X is

$$\mathcal{E} = \{v \mid v^\top \mathrm{diag}((\delta M - \varepsilon)^2, M^2)v = (\delta M - \varepsilon)^2 v_1^2 + M^2 v_2^2 \le M^2(\delta M - \varepsilon)^2\}.$$

The third column of X is outside this ellipsoid as can be seen in Figure 5, since

$$\frac{4\varepsilon^2(\delta M - \varepsilon)^2/\delta^2 + (\delta M - \varepsilon)^2 M^2}{M^2(\delta M - \varepsilon)^2} > 1.$$

As a consequence, the minimum volume ellipsoid containing all column s of X must have the third column on its boundary, and MVE-SPA may identify it after the preconditioning, since all the columns of the preconditioned matrix $\widetilde{X}$ have norm 1 as can be seen in Figure 6. The distance between the third column and the first is at least

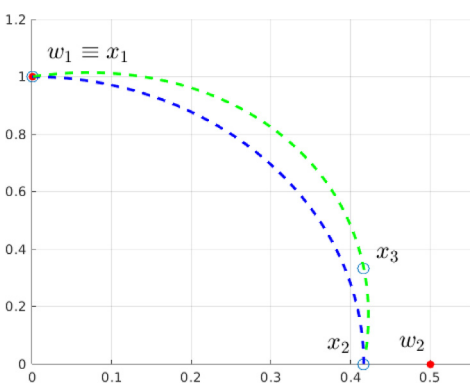

FIG. 5. *Columns of matrix X as empty blue circles and columns of matrix W as red dots. The dashed blue line is the minimum volume ellipsoid containing x_1 and x_2. The dashed green line is the minimum volume ellipsoid containing all columns of X.*

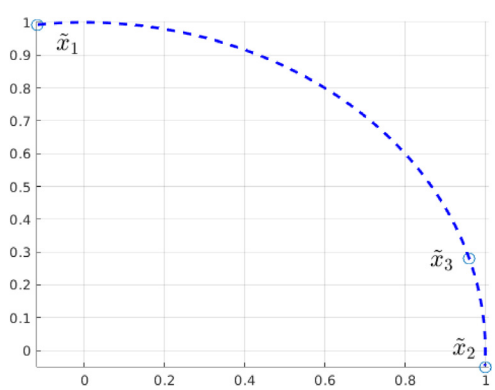

FIG. 6. *Columns of the preconditioned matrix. The points on the circle (dashed blue line) have norm* 1.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$$M - \frac{2\varepsilon}{\delta} > \frac{2\varepsilon}{\delta} = 2\mathcal{K}(W)\varepsilon,$$

and the distance between the third column and the second is exactly $\frac{2\varepsilon}{\delta} = 2\mathcal{K}(W)\varepsilon$. This concludes that the error bound of MVE-SPA will be larger than $2\mathcal{K}(W)\varepsilon$. ∎

7. Translation and lifting. We have seen that SPA cannot extract r vertices in dimension $r-1$. This can be resolved using T-SPA (see Algorithm 2), via a translation instead of a projection at the first step, given that $H^\top e = e$, as discussed in section 2.1. As briefly mentioned in section 2.1, this is also possible by lifting the data points in one dimension higher adding a vector $ce^\top$ for $c \in \mathbb{R}$ as the last row of X (see (2.1)); see also, e.g., [27].

In this section, we propose a new strategy, combining the lifting procedure with a translation, in order to minimize the conditioning of W. In fact, under the condition that $H^\top e = e$, translating and shifting does not change the solution H and translates and shifts W in the same way since

$$X = WH \iff \begin{pmatrix} X - ve^\top \\ ce^\top \end{pmatrix} = \begin{pmatrix} W - ve^\top \\ ce^\top \end{pmatrix} H \tag{7.1}$$

for any vector v and constant c. The problem is as follows: find $v \in \mathbb{R}^m$ and $c \in \mathbb{R}$ such that the conditioning of $\left(\begin{smallmatrix} W - ve^\top \\ ce^\top \end{smallmatrix}\right)$ is minimized. Of course, W is unknown.

In the following, we propose a simple heuristic to tackle this problem. In order to decrease the largest singular value of $W - ve^\top$, a heuristic is to minimize the sum of its squared singular values, that is, $\|W - ve^\top\|_F^2$. The optimal solution of this problem is given by the average of the column of W, that is, $v = \frac{1}{r}We$. Note that if it happens the best rank-one approximation of W has the form $ue^\top$ for some vector u (in other terms, $e/\sqrt{r}$ is the first right singular vector of W), and then the largest singular value of $W - ue^\top$ becomes $\sigma_2(W)$, which coincides with the smallest $\sigma_1(W - ze^\top)$ among all z (by the interlacing property of singular values applied to rank 1 modifications). In practice, this average is typically well approximated by the average of the columns of X, given that the data points are evenly distributed in $\operatorname{conv}(W)$. Hence, we suggest to use $v = \frac{1}{n}Xe$. It is also possible to use more advanced strategies, e.g., take v as the average of the columns of X extracted by SPA [1]. However, for simplicity, we consider in this paper only $v = \frac{1}{n}Xe$. Then, assuming $v = \frac{1}{r}We$, we can characterize the conditioning of $\left(\begin{smallmatrix} W - ve^\top \\ ce^\top \end{smallmatrix}\right)$ as a function of c as follows, from which we will propose a value for c.

LEMMA 5. *Suppose $W_t \in \mathbb{R}^{m \times r}$ with $m \geq r-1$ and the average of its columns is equal to 0, that is, $W_t e = 0$. If $W_\ell = \left(\begin{smallmatrix} W_t \\ c\, e^\top \end{smallmatrix}\right)$, then*

$$\mathcal{K}(W_\ell) = \begin{cases} c\sqrt{r}/\sigma_{r-1}(W_t) & \text{if } c \geq \|W_t\|/\sqrt{r}, \\ \|W_t\|/\sigma_{r-1}(W_t) & \text{if } \|W_t\|/\sqrt{r} \geq c \geq \sigma_{r-1}(W_t)/\sqrt{r}, \\ \|W_t\|/(c\sqrt{r}) & \text{if } \sigma_{r-1}(W_t)/\sqrt{r} \geq c. \end{cases} \tag{7.2}$$

Proof. If we decompose $x \in \mathbb{R}^r$ as $x = \alpha e/\sqrt{r} + w\sqrt{1-\alpha^2}$, where $\|w\|^2 = 1$ and w is orthogonal to e, then

$$\begin{aligned} \|W_\ell\|^2 &= \max_{\|x\|=1} \|W_\ell x\|^2 = \max_{\|x\|=1} \|W_t x\|^2 + (ce^\top x)^2 \\ &= \max_{\|w\|=1, w^\top e=0, |\alpha|\leq 1} (1-\alpha^2)\|W_t w\|^2 + (\alpha c)^2 r \\ &= \max\left\{ \max_{\|w\|=1, w^\top e=0} \|W_t w\|^2, c^2 r \right\} = \max\left\{ \|W_t\|^2, c^2 r \right\}, \end{aligned}$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Algorithm 3. Translated lifted SPA.

Input: The matrix $X = WH + N$, where $H \geq 0$, $H^\top e = e$, H is separable, W is full column rank, the number r of columns to extract.

Output: Set of r indices $\mathcal{J}$ such that $X(:,\mathcal{J}) \approx W$ (up to permutation).

1: $v = Xe/n$.
2: $c = \frac{1}{2}(\min_i \|X(:,i) - v\|/\sqrt{r} + \|X - ve^\top\|_F/\sqrt{n})$.
3: $X_{t\ell} = \left(\begin{smallmatrix} X - ve^\top \\ c\,e^\top \end{smallmatrix}\right)$. *% Translate and lift*
4: $\mathcal{J} = \mathrm{SPA}(X_{t\ell}, r)$. *% see Algorithm 1*

so $\|W_\ell\| = \max\{\|W_t\|, c\sqrt{r}\}$. In a totally analogous way we find that $\sigma_r(W_\ell) = \min\{\sigma_{r-1}(W_t), c\sqrt{r}\}$. ∎

According to (7.2), assuming $v = We/r$, and denoting by $W_t = W - ve^\top$ the translated version of W, and by $W_\ell = \left(\begin{smallmatrix} W_t \\ c\,e^\top \end{smallmatrix}\right)$ the lifted version of W, the value of c minimizing $\mathcal{K}(W_\ell)$ is in the interval

$$\left[\frac{1}{\sqrt{r}}\sigma_{r-1}\left(W - ve^\top\right), \frac{1}{\sqrt{r}}\left\|W - ve^\top\right\|\right].$$

How can we approximate this interval since W is unknown? Since the columns of $X_t = X - ve^\top$ are within the convex hull of the columns of $W_t = W - ve^\top$, we assume $\|W_t\|_F \approx \frac{\sqrt{r}}{\sqrt{n}}\|X_t\|_F$, and hence we approximate the upper bound of the interval using the following approximation: $\|W_t\|_F/\sqrt{r} \approx \|X_t\|_F/\sqrt{n}$. For the lower bound, we estimate $\sigma_{r-1}(W_t)$ with $\min_j \|X_t(:,j)\|$. Finally, we propose to use c as the average of the estimated lower and upper bounds,

$$c = \frac{1}{2}\left(\min_i \|X(:,i) - v\|/\sqrt{r} + \left\|X - ve^\top\right\|_F/\sqrt{n}\right). \tag{7.3}$$

However, we have not observed a significant sensitivity of TL-SPA w.r.t. to the value of c, as hinted in Lemma 5. For example, using the average of the absolute values of $X - ve^\top$ provides similar results.

We will refer to the variant of SPA that applies SPA on W_ℓ, the translated and lifted data, as translated lifted SPA (TL-SPA); see Algorithm 3.

Remark 4 (implementation of TL-SPA). As for SPA and T-SPA (see Remarks 1 and 2), TL-SPA can be implemented in $\mathcal{O}(r\,\mathrm{nnz}(X))$ operations (essentially, r matrix-vector products). In particular, $X - ve^\top$ should not be computed explicitely as v could be dense. See https://gitlab.com/ngillis/robustSPA for a MATLAB implementation.

8. Numerical experiments. In this section, we illustrate some of the theoretical findings by comparing the various algorithms discussed in this paper, namely, the following:

- SPA [4, 19]; see Algorithm 1.
- T-SPA [4]; see Algorithm 2.
- FAW [4]: it starts with T-SPA and then loops once over all extracted by T-SPA and tries to replace them by increasing the volume of the convex hull of the extracted columns of X; see also section 2.2.
- SPA^2 [18]: SPA preconditioned with SPA; see section 6.1.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

- TL-SPA: it applies SPA on the translated and lifted data matrix as proposed in section 7.
- TL-SPA^2: it applies SPA^2 on the translated and lifted data matrix as described in section 7.

We do not include MVE-SPA as it performs similarly as SPA^2, as reported, e.g., in [15, Chapter 7]. The code used to run these experiments is available from https://gitlab.com/ngillis/robustSPA. We do not report computational times as all algorithms are extremely fast, although FAW is about r times slower than SPA, T-SPA, and TL-SPA, while SPA^2 and TL-SPA^2 are about two times slower. (It takes about 10 s on a standard laptop to run the experiments described below, running the 6 algorithms on 6120 matrices.)

Synthetic data sets. We use the same synthetic data sets as in [15, Chapter 7.4.6.1]: $X = WH + N$, where the following hold:

- The entries of $W \in \mathbb{R}^{m\times r}$ are generated uniformly at random in the interval $[0,1]$, `W = rand(m,r)`. In this case, W is well-conditioned with high probability. To make W ill-conditioned, we compute the compact SVD of $W = U\Sigma V^\top$, then update $W \leftarrow U\Sigma' V^\top$, where Σ' is a diagonal matrix whose entries are log-spaced between 10^{-6} and 1, and hence $\frac{\sigma_1(W)}{\sigma_r(W)} = 10^6$ and $\mathcal{K}(W)$ will be smaller, of the order of $5 \cdot 10^5$; see Table 3.
- The matrix $H = [I_r, H']\Pi \in \mathbb{R}^{r\times n}$, where Π is a random permutation, and the columns of H' either follow the Dirichlet distribution of parameters 0.5, or each column of H' has two nonzero entries equal to 1/2 for all possible combinations of two columns of W, so that $n = r + \binom{r}{2}$. The latter case is the so-called middle point experiment since besides the columns of W, the columns of X are in-between two columns of W. This is a more adversarial case than the Dirichlet distribution, as all data points are on edges of the boundary of $\text{conv}(W)$.
- In the case of the Dirichlet distribution, each entry of N follows a Gaussian distribution, `N = randn(m,n)`, and is then scaled $N \leftarrow \delta \frac{\|WH\|_F}{\|N\|_F} N$, where δ is the noise level (=ratio of the norms of N and WH). For the middle point experiments, the middle points are moved toward the outside of the convex hull of $\text{conv}(W)$: if $X(:,j)$ is a middle point, $N(:,j) = \delta(X(:,j) - \bar{w})$, where $\bar{w}$ is the average of the columns of W. The columns of W in X are not perturbed. This is an adversarial noise: the data points are moved outside $\text{conv}(W)$.

Table 1 summarizes the four types of data sets we will generate. For each experimental setting, we test 51 different noise levels δ and, for each noise level, we report the average accuracy (which is the proportion of the indices of the columns of W correctly recovered by the corresponding algorithm) over 30 randomly generated matrices.

Figure 7 reports the accuracy as a function of the noise level for the Dirichlet experiments (Exp. 1 and 2), and Figure 8 for the middle point experiments (Exp. 3 and 4).

Table 1
Four types of synthetic data sets used in our four experiments.

	m	n	r	W	H	N
Exp. 1	40	110	10	Well-cond.	Dirichlet	Gaussian
Exp. 2	40	110	10	Ill-cond.	Dirichlet	Gaussian
Exp. 3	40	55	10	Well-cond.	Middle points	adversarial
Exp. 4	9	55	10	Rank-def.	Middle points	adversarial

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

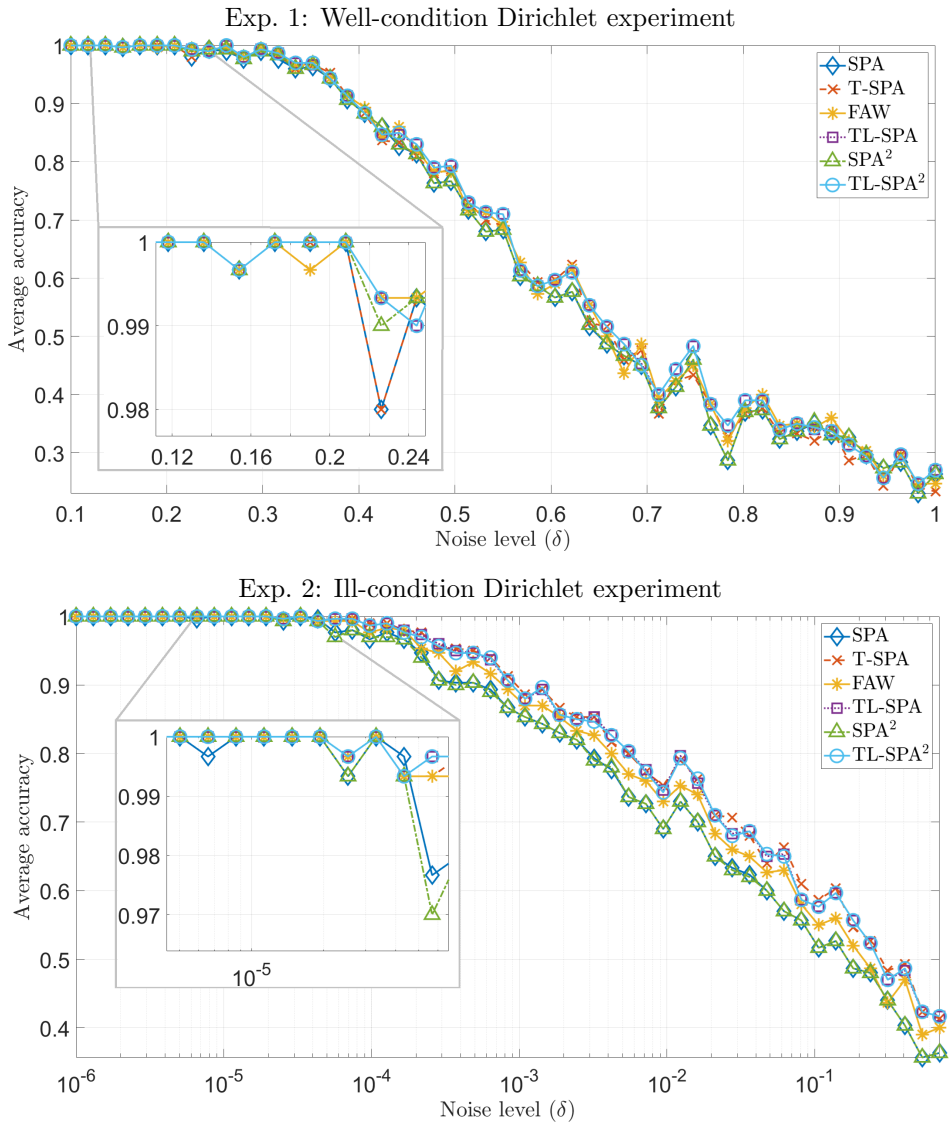

FIG. 7. *Accuracy vs. noise level for Dirichlet experiments.*

Table 2 reports the robustness of each algorithm in each experimental setting. The robustness is defined as the largest noise level for which the corresponding algorithm recovered exactly all indices for all matrices.

We observe the following:

- For the well-conditioned Dirichlet experiment (Exp. 1), all SPA variants perform similarly and have the same robustness. The main reason is that the conditioning of W and its translated variant, W_t, are close to each other. This is a similar observation as in [15, Figure 7.6].
- For the ill-conditioned Dirichlet experiment (Exp. 2), SPA and SPA^2 perform slightly worse than the other algorithms. The reason is that W has a very small eigenvalue (10^{-6}) that is slightly better handled via translation at the first step. However, the gain of T-SPA, FAW, TL-SPA, and TL-SPA^2

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

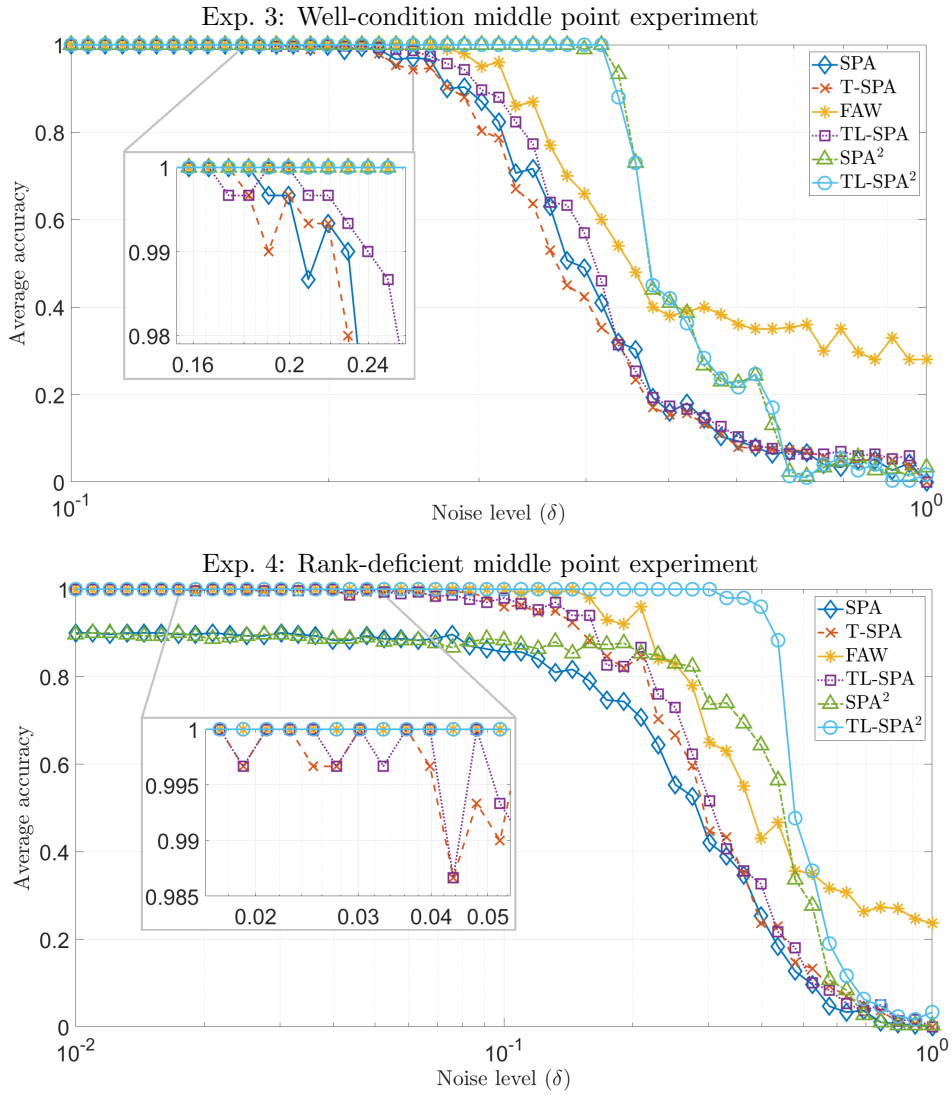

FIG. 8. *Accuracy vs. noise level for the middle point experiments.*

is not significant because the second singular value is also small ($=4.610^{-6}$ due to the log-spaced singular values) so that the conditioning of W_t is not significantly smaller than that of W.

- For the well-conditioned middle point experiment (Exp. 3), there is a clear hierarchy: SPA and T-SPA perform similarly (since the matrix is well-conditioned, translation does not bring much improvement), FAW is more robust than TL-SPA but its accuracy decays more rapidly, SPA^2 outperforms FAW (as already noted in [15, Figure 7.7]), while TL-SPA^2 and SPA^2 perform similarly although TL-SPA^2 is more robust (see Table 2).
- For the rank-deficient middle point experiment (Exp. 4), SPA and SPA^2 cannot extract more than nine indices, since $m = 9 < r = 10$, and hence cannot have an accuracy larger than 90%. This is because W has only $r-1$ rows and hence as an infinite conditioning (the rth singular value is equal to

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

TABLE 2

Robustness comparison of SPA-like algorithms: largest value of the noise level δ for which each algorithm recovers all indices correctly for the 30 generated matrices. The best result is highlighted in bold.

	SPA	T-SPA	FAW	TL-SPA	SPA^2	TL-SPA^2
Exp. 1	**0.136**	**0.136**	**0.136**	**0.136**	**0.136**	**0.136**
Exp. 2	$5.03\ 10^{-6}$	$\mathbf{1.93\ 10^{-5}}$	$\mathbf{1.93\ 10^{-5}}$	$\mathbf{1.93\ 10^{-5}}$	$\mathbf{1.93\ 10^{-5}}$	$\mathbf{1.93\ 10^{-5}}$
Exp. 3	0.182	0.174	0.263	0.166	0.380	**0.417**
Exp. 4	0	0.017	0.100	0.017	0	**0.302**

TABLE 3

Average condition number, and standard deviation, of the matrix W after preprocessing using the different algorithms, for the 30 generated matrices (same settings as in Table 2). The first column reports the noise level δ. The best result is highlighted in bold.

Exp. 1	SPA	T-SPA & FAW	TL-SPA	SPA^2	TL-SPA^2
0.10	3.7 ± 0.4	2.7 ± 0.3	4.5 ± 1.4	$\mathbf{1.2} \pm 0.0$	$\mathbf{1.2} \pm 0.0$
0.28	3.9 ± 0.4	2.8 ± 0.4	1.9 ± 0.2	1.8 ± 0.2	$\mathbf{1.7} \pm 0.2$
0.46	3.7 ± 0.4	2.9 ± 0.4	$\mathbf{2.0} \pm 0.2$	23.6 ± 25.8	22.6 ± 52.8
0.64	3.7 ± 0.5	3.1 ± 0.4	$\mathbf{2.1} \pm 0.2$	75.4 ± 241.1	212.4 ± 775.5
0.82	3.8 ± 0.3	3.6 ± 0.4	$\mathbf{2.5} \pm 0.2$	58.1 ± 78.8	135.5 ± 501.8
1	3.8 ± 0.3	4.0 ± 0.5	$\mathbf{2.6} \pm 0.2$	49.1 ± 60.1	283.6 ± 621.0
Exp. 2	**SPA**	**T-SPA & FAW**	**TL-SPA**	**SPA^2**	**TL-SPA^2**
10^{-6}	$54.7 \pm 5.55 \cdot 10^4$	$20.4 \pm 1.93 \cdot 10^4$	$11.9 \pm 1.20 \cdot 10^4$	1.1 ± 0.0	$\mathbf{1.0} \pm 0.0$
$1.48\ 10^{-5}$	$55.9 \pm 7.69 \cdot 10^4$	$19.6 \pm 2.03 \cdot 10^4$	$12.2 \pm 1.65 \cdot 10^4$	6.6 ± 2.8	$\mathbf{1.3} \pm 0.1$
$2.18\ 10^{-4}$	$59.4 \pm 7.68 \cdot 10^4$	$51.2 \pm 5.66 \cdot 10^3$	$13.0 \pm 1.75 \cdot 10^4$	802.3 ± 1157.2	$\mathbf{118.1} \pm 122.7$
0.003	$58.5 \pm 7.81 \cdot 10^4$	$48.4 \pm 4.63 \cdot 10^3$	$12.8 \pm 1.74 \cdot 10^4$	$1.13 \pm 1.51 \cdot 10^4$	$\mathbf{4.57} \pm 8.37 \cdot 10^3$
0.047	$58.5 \pm 7.39 \cdot 10^4$	$49.6 \pm 4.59 \cdot 10^3$	$12.7 \pm 1.56 \cdot 10^4$	$8.01 \pm 30.9 \cdot 10^5$	$\mathbf{2.19} \pm 3.13 \cdot 10^4$
0.7	$57.4 \pm 6.06 \cdot 10^4$	$\mathbf{49.1} \pm 4.68 \cdot 10^3$	$12.7 \pm 1.26 \cdot 10^4$	$2.21 \pm 5.67 \cdot 10^6$	$4.31 \pm 5.94 \cdot 10^5$
Exp. 3	**SPA**	**T-SPA & FAW**	**TL-SPA**	**SPA^2**	**TL-SPA^2**
0.10	3.9 ± 0.4	2.7 ± 0.3	2.2 ± 0.2	$\mathbf{1.0} \pm 0.0$	$\mathbf{1.0} \pm 0.0$
0.16	3.7 ± 0.4	2.6 ± 0.3	2.2 ± 0.2	$\mathbf{1.0} \pm 0.0$	$\mathbf{1.0} \pm 0.0$
0.25	3.7 ± 0.4	2.6 ± 0.3	2.2 ± 0.2	1.6 ± 0.7	$\mathbf{1.2} \pm 0.5$
0.40	3.7 ± 0.4	2.6 ± 0.3	$\mathbf{2.2} \pm 0.2$	2.7 ± 0.2	2.7 ± 0.2
0.63	3.8 ± 0.6	3.0 ± 0.4	$\mathbf{2.3} \pm 0.3$	3.2 ± 0.4	3.1 ± 0.3
1	3.9 ± 0.4	3.3 ± 0.4	$\mathbf{2.3} \pm 0.2$	3.5 ± 0.4	3.6 ± 0.4
Exp. 4	**SPA**	**T-SPA & FAW**	**TL-SPA**	**SPA^2**	**TL-SPA^2**
0.01	∞	72.0 ± 95.5	57.7 ± 73.8	∞	$\mathbf{1.0} \pm 0.0$
0.025	∞	53.1 ± 72.0	42.6 ± 57.6	∞	$\mathbf{1.0} \pm 0.0$
0.063	∞	$7.65 \pm 38.3 \cdot 10^3$	$6.51 \pm 32.8 \cdot 10^3$	∞	$\mathbf{1.2} \pm 0.5$
0.16	∞	407 ± 1610	324 ± 1252	∞	$\mathbf{1.8} \pm 0.7$
0.40	∞	56.5 ± 94.6	41.4 ± 72.1	∞	$\mathbf{2.9} \pm 0.4$
1	∞	195 ± 600	129 ± 416	∞	$\mathbf{3.5} \pm 0.4$

zero), while W_t is well conditioned (since the $(r-1)$th singular value of W is not small, with high probability). For the other algorithms, we observe a similar hierarchy as for the well-conditioned middle point experiment (Exp. 3), except that our newly proposed TL-SPA^2 significantly outperforms the other variants (robustness of 0.302 vs. 0.100 for the second best, FAW).

These experiments illustrate that, in adversarial and rank-deficient or ill-conditioned settings, it is key to use enhanced variants of SPA. In particular, the proposed TL-SPA^2 has been shown to outperform all other variants in such cases.

To further validate this observation, Table 3 reports the conditioning of the matrix W after the preprocessing performed by the different algorithms, namely, the following hold:

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

TABLE 4
Average conditioning $\mathcal{K}(W_\ell)$ using different shifts c defined in (7.3) multiplied by a factor $\alpha \in \{0.1, 0.5, 2, 10\}$. The same settings as in Table 3 but the results are averaged over all noise levels. The best result is highlighted in bold.

	$\alpha = 1/10$	$\alpha = 1/2$	$\alpha = 1$	$\alpha = 2$	$\alpha = 10$
Exp. 1	6.3 ± 6.4	3.0 ± 2.6	$\mathbf{2.6} \pm 1.1$	3.0 ± 0.9	10.6 ± 6.2
Exp. 2	$\mathbf{12.4} \pm 1.60 \cdot 10^4$	$\mathbf{12.4} \pm 1.58 \cdot 10^4$	$12.5 \pm 1.57 \cdot 10^4$	$13.0 \pm 1.54 \cdot 10^4$	$22.9 \pm 2.44 \cdot 10^4$
Exp. 3	4.2 ± 0.3	$\mathbf{1.9} \pm 0.2$	2.2 ± 0.2	3.2 ± 0.4	13.6 ± 1.6
Exp. 4	$\mathbf{1.03} \pm 11.8 \cdot 10^3$	$1.07 \pm 12.2 \cdot 10^3$	$1.18 \pm 13.4 \cdot 10^3$	$1.56 \pm 17.6 \cdot 10^3$	$5.92 \pm 66.4 \cdot 10^3$

- For SPA: $\mathcal{K}(W)$, since SPA applies directly to the data.
- For T-SPA and FAW: $\frac{K(W_t)}{\sigma_{r-1}(W_t)}$, where W_t is the matrix W translated by the column of X with the largest ℓ_2 norm; see section 2.1.
- For TL-SPA: $\mathcal{K}(W_\ell)$, where $W_\ell = \left(\begin{smallmatrix} W - ve^\top \\ c\, e^\top \end{smallmatrix}\right)$ with $v = Xe/n$ and c given by (7.3).
- For SPA^2: $\mathcal{K}((X(:,\mathcal{J}))^\dagger W)$, where $\mathcal{J}$ is obtained by SPA.
- For TL-SPA^2: $\mathcal{K}((\tilde{X}(:,\mathcal{J}))^\dagger W_\ell)$, where $\mathcal{J}$ is obtained by TL-SPA, and $\tilde{X} = \left(\begin{smallmatrix} X - ve^\top \\ c\, e^\top \end{smallmatrix}\right)$ is the translated and lifted version of X.

We observe the following:

- Except for Exp. 2 and $\delta = 0.7$, it is always TL-SPA or TL-SPA^2 that lead to the lowest condition numbers.
- For Exp. 1, where W is well-conditioned, the variants using the left-inverse preconditioning (that is, TL-SPA^2, and to a lesser extent SPA^2) lead to high condition numbers for high noise levels (for $\delta \geq 0.46$). However, this is not really worrisome since at these noise levels, the performance of all algorithms degrades significantly (robustness is lost for $\delta < 0.136$; see Table 2).
- Exp. 2, where W is ill-conditioned, TL-SPA^2 is very effective at reducing the conditioning, especially for low noise levels, which explains its robustness results from Table 2.
- Exp. 4, where W is rank-deficient, TL-SPA^2 is very effective in reducing the conditioning, when compared to T-SPA and TL-SPA, again confirming the results from Table 2.

Remark 5 (sensitivity of TL-SPA to the lift parameter). In section 7, we proposed a lift parameter for TL-SPA based on Lemma 5; see (7.3). This is a heuristic choice with no theoretical guarantees since W is unknown in practice. However, we have observed that TL-SPA is not too sensitive to the choice of this parameter. We have rerun TL-SPA with the lift parameter c defined in (7.3) multiplied by a factor $\alpha \in \{0.1, 0.5, 2, 10\}$. Table 4 reports the average condition number, $\mathcal{K}(W_\ell)$, for all experiments. Except when $\alpha = 10$, which leads to significantly worse results, the other condition numbers are very close to each other.

9. Conclusion and further work. In this paper, we revisited the robustness to noise of SPA and variants. We improved previous error bounds for the first step of SPA and when $r = 2$ (Theorems 3 and 4) and for the first two steps and when $r = 3$ for the translated variant of SPA, T-SPA (Theorems 5 and 6). We then proved that, for $r \geq 3$, the original robustness result of SPA [19, Theorem 3] is tight (Theorem 7), an important open question in the literature [15, p. 231], using a well-constructed family of matrices that achieve the worst-case error bound. We also proved the tightness of the error bounds for two preconditioned variants of SPA, namely, SPA preconditioned

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

with itself (SPA^2, Theorem 9) and SPA preconditioned with MVE (MVE-SPA, Theorem 11). Finally, we proposed a new SPA variant that uses a translation and lifting of the data points, namely, TL-SPA, and that allows us to improve the robustness of SPA by reducing the condition number of W. We illustrated this on numerical experiments.

Further work includes the following:

- TL-SPA does not have provable robustness guarantees as the choice of the translation and lifting may influence the conditioning of W in an unexpected way. Providing such guarantees would be important to better understand its behavior. This could also lead to new, more clever, ways to preprocess the input data matrix X.
- Tightness of other separable SSMF algorithms could also be studied, e.g., the ones based on convex relaxations (see section 2.2). Maybe our counterexamples for SPA, SPA^2, and MVE-SPA could be the basis for such an analysis.
- Comparing the performance of the SPA variants for real-world applications would help understand which variant to use in which situation.

Appendix A. Proofs.

A.1. Proof of Lemma 1.

Proof. Given a couple of distinct indices i, j, let $W_{i,j} = [w_i, w_j]$. We have

$$\begin{aligned}\left\|\frac{w_i+w_j}{2}\right\|^2 &= \frac{1}{2}\|w_i\|^2 + \frac{1}{2}\|w_j\|^2 - \frac{1}{4}\|w_i - w_j\|^2 \le K(W_{i,j})^2 - \frac{1}{4}\|W_{i,j}(e_1 - e_2)\|^2 \\ &\le K(W_{i,j})^2 - \frac{1}{2}\sigma_2(W_{i,j})^2.\end{aligned}$$

This implies

$$\begin{aligned}\left\|\frac{w_i+w_j}{2}\right\| &\le \sqrt{K(W_{i,j})^2 - \frac{1}{2}\sigma_2(W_{i,j})^2} \le \sqrt{K(W_{i,j})^2 - \frac{1}{2}\sigma_2(W_{i,j})^2 + \frac{\sigma_2(W_{i,j})^4}{16K(W_{i,j})^2}} \\ &= K(W_{i,j}) - \frac{\sigma_2(W_{i,j})^2}{4K(W_{i,j})} = K(W_{i,j}) - \frac{\sigma_2(W_{i,j})}{4\mathcal{K}(W_{i,j})}.\end{aligned}$$

By the theorem of interlacing of singular values $\sigma_2(W_{i,j}) \ge \sigma_r(W)$ and $\mathcal{K}(W_{i,j}) \le \mathcal{K}(W)$, so

$$\begin{aligned}\left\|\frac{w_i+w_j}{2}\right\| &\le K(W_{i,j}) - \frac{\sigma_2(W_{i,j})}{4\mathcal{K}(W_{i,j})} \le K(W_{i,j}) - \frac{\sigma_r(W)}{4\mathcal{K}(W)} \\ &\le \max\{\|w_i\|, \|w_j\|\} - 2\varepsilon.\end{aligned}$$

∎

A.2. Proof of Lemma 2.

Proof of Lemma 2. Notice that since all w_i and all v_i with $i \ge 1$ are in C, then $\operatorname{conv}(v_1, \ldots, v_r, w_1, \ldots, w_r) \subseteq C$. As a consequence, to prove $V \backslash C \subseteq \operatorname{conv}(v_0, w_1, \ldots, w_r)$ we just need to show that

$$V = \operatorname{conv}(v_1, \ldots, v_r, w_1, \ldots, w_r) \cup \operatorname{conv}(v_0, w_1, \ldots, w_r).$$

Let $\widetilde{v}_i := v_i - v_0$ and $\widetilde{w}_i := w_i - v_0 = \alpha_i \widetilde{v}_i$. By the definition of convex hull, $v \in V$ if and only if there exist nonnegative $\lambda_0, \ldots, \lambda_r$ that sum to one such that $\widetilde{v} := v - v_0 = -v_0 + \sum_i \lambda_i v_i = \sum_{i \ge 1} \lambda_i \widetilde{v}_i$. If $\sum_{i \ge 1} \lambda_i / \alpha_i \le 1$, then

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

$$\widetilde{v} = \sum_{i\geq 1} \frac{\lambda_i}{\alpha_i}\widetilde{w}_i \implies v = v_0\left(1 - \sum_{i\geq 1}\frac{\lambda_i}{\alpha_i}\right) + \sum_{i\geq 1}\frac{\lambda_i}{\alpha_i} w_i \in \operatorname{conv}(v_0, w_1, \ldots, w_r).$$

Otherwise, $\sum_{i\geq 1}\lambda_i/\alpha_i > 1$, but since $\sum_{i\geq 1}\lambda_i = 1-\lambda_0 \leq 1$, then there exists $t\in[0,1)$ such that $\sum_{i\geq 1}\lambda_i/(1+t(\alpha_i-1)) = 1$. Let now $\widetilde{\lambda}_i := \lambda_i/(1+t(\alpha_i-1))$ for every $i\geq 1$, so that $\sum_{i\geq 1}\widetilde{\lambda}_i = 1$ and

$$t\widetilde{\lambda}_i\widetilde{w}_i + (1-t)\widetilde{\lambda}_i\widetilde{v}_i = (1+(\alpha_i-1)t)\widetilde{\lambda}_i\widetilde{v}_i = \lambda_i\widetilde{v}_i.$$

Therefore

$$\begin{aligned}\widetilde{v} &= \sum_{i\geq 1} t\widetilde{\lambda}_i\widetilde{w}_i + (1-t)\widetilde{\lambda}_i\widetilde{v}_i \implies v\\ &= \sum_{i\geq 1} t\widetilde{\lambda}_i w_i + (1-t)\widetilde{\lambda}_i v_i \in \operatorname{conv}(v_1,\ldots,v_r,w_1,\ldots,w_r).\end{aligned}$$

∎

A.3. Proof of Lemma 3.

Proof. Notice that $\sigma_{r-1}(\hat{W}) = \sigma_{r-1}(\widetilde{W} - ve^\top)$, where

$$\widetilde{W} = (w_1, w_2, \ldots, w_{r-1}, v) = W\begin{pmatrix} & \lambda_1\\ I_{r-1} & \vdots\\ & \lambda_{r-1}\\ 0_{1\times(r-1)} & \lambda_r\end{pmatrix} = W(I + \widetilde{\lambda}e_r^\top)$$

with $\widetilde{\lambda}_i = \lambda_i$ for $1\leq i<r$ and $-\widetilde{\lambda}_r = 1-\lambda_r = \lambda_1+\cdots+\lambda_{r-1}$. Due to the interlacing property of singular values, we have that $\sigma_{r-1}(\widetilde{W}-ve^\top) \geq \sigma_r(\widetilde{W}) \geq \sigma_r(W)\sigma_r(I+\widetilde{\lambda}e_r^\top)$. Moreover, due to Weyl's perturbation law,

$$\begin{aligned}\sigma_r(I+\widetilde{\lambda}e_r^\top) &\geq 1 - \|\widetilde{\lambda}\| = 1 - \sqrt{\lambda_1^2+\cdots+\lambda_{r-1}^2+(1-\lambda_r)^2}\\ &\geq 1 - \sqrt{2(1-\lambda_r)^2} \geq 1 - \frac{\sqrt{2}}{2}.\end{aligned}$$

This lets us conclude that $\sigma_{r-1}(\hat{W}) \geq \sigma_r(W)\frac{2-\sqrt{2}}{2}$. For the second part, it is enough to notice that $\|v\| \leq K(W)$ since it belongs to $\operatorname{conv}(W)$, so $K(\hat{W}) \leq 2K(W)$ is immediate. ∎

A.4. Proof of Lemma 4.

Proof. Since $c\leq 1/8$, we can apply Theorem 3 and find that $\|x_1-w_1\| \leq 33\mathcal{K}(W)\varepsilon$. The translation brings the matrix X into $\hat{X}$ as in the hypothesis. From Lemma 3, we get that

$$\sigma_{r-1}(\hat{W}) \geq \frac{2-\sqrt{2}}{2}\sigma_r(W) > 0,$$

so $\hat{W}\in\mathbb{R}^{m\times(r-1)}$ is full rank. Notice also that

$$\begin{aligned}\|r_1\| &\leq \|x_1-w_1\| + \|x_1-w\| \leq 34\mathcal{K}(W)\varepsilon \leq 34c\sigma_r(W)\\ &\leq \frac{68c}{2-\sqrt{2}}\sigma_{r-1}(\hat{W}) \leq \frac{1}{2}\sigma_{r-1}(\hat{W}).\end{aligned}$$

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

Last, $K(\hat{W}) \leq 2K(W)$, which implies

$$\sigma_r(W)\mathcal{K}(W)^{-1} = \sigma_r(W)^2 K(W)^{-1} \leq \frac{8}{(2-\sqrt{2})^2}\sigma_{r-1}(\hat{W})^2 K(\hat{W})^{-1}$$
$$= \frac{4}{3-2\sqrt{2}}\sigma_{r-1}(\hat{W})\mathcal{K}(\hat{W})^{-1}$$

and

$$\mathcal{K}(\hat{W}) = \sigma_{r-1}(\hat{W})^{-1}K(\hat{W}) \leq \frac{4}{2-\sqrt{2}}\sigma_r(W)^{-1}K(W) = \frac{4}{2-\sqrt{2}}\mathcal{K}(W). \qquad \square$$

Reproducibility of computational results. This paper has been awarded the "SIAM Reproducibility Badge: Code and data available" as a recognition that the authors have followed reproducibility principles valued by SIMAX and the scientific computing community. Code and data that allow readers to reproduce the results in this paper are available at https://gitlab.com/ngillis/robustspa.

Acknowledgments. We are grateful to the anonymous reviewers who carefully read the manuscript; their feedback helped us improve our paper. The first author is a member of the Research Group GNCS (Gruppo Nazionale per il Calcolo Scientifico) of INdAM (Istituto Nazionale di Alta Matematica).

REFERENCES

[1] M. Abdolali and N. Gillis, *Simplex-structured matrix factorization: Sparsity-based identifiability and provably correct algorithms*, SIAM J. Math. Data Sci., 3 (2021), pp. 593–623, https://doi.org/10.1137/20M1354982.

[2] J. Agterberg and A. R. Zhang, *Estimating higher-order mixed memberships via the $\ell_{2,\infty}$ tensor perturbation bound*, J. Amer. Statist. Assoc., 120 (2025), pp. 1214–1224, https://doi.org/10.1080/01621459.2024.2404265.

[3] U. Araújo, B. Saldanha, R. Galvão, T. Yoneyama, H. Chame, and V. Visani, *The successive projections algorithm for variable selection in spectroscopic multicomponent analysis*, Chemometr. Intell. Lab. Syst., 57 (2001), pp. 65–73, https://doi.org/10.1016/S0169-7439(01)00119-8.

[4] S. Arora, R. Ge, Y. Halpern, D. Mimno, A. Moitra, D. Sontag, Y. Wu, and M. Zhu, *A practical algorithm for topic modeling with provable guarantees*, in Proceedings of the International Conference on Machine Learning, 2013, pp. 280–288.

[5] A. Bakshi, C. Bhattacharyya, R. Kannan, D. P. Woodruff, and S. Zhou, *Learning a latent simplex in input sparsity time*, in Proceedings of the International Conference on Learning Representations (ICLR), 2021.

[6] C. Bhattacharyya and R. Kannan, *Finding a latent k–simplex in $o^*(k \cdot nnz(data))$ time via subset smoothing*, in Proceedings of the ACM-SIAM Symposium on Discrete Algorithms, SIAM, 2020, pp. 122–140.

[7] O. S. Borgen and B. R. Kowalski, *An extension of the multivariate component-resolution method to three components*, Anal. Chim. Acta, 174 (1985), pp. 1–26, https://doi.org/10.1016/S0003-2670(00)84361-5.

[8] P. Businger and G. Golub, *Linear least squares solutions by Householder transformations*, Numer. Math., 7 (1965), pp. 269–276, https://doi.org/10.1007/BF01436084.

[9] E. J. Candès, X. Li, Y. Ma, and J. Wright, *Robust principal component analysis?*, J. ACM, 58 (2011), 11, https://doi.org/10.1145/1970392.1970395.

[10] E. Elhamifar, G. Sapiro, and R. Vidal, *See all by looking at a few: Sparse modeling for finding representative objects*, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, IEEE, 2012, pp. 1600–1607.

[11] E. Esser, M. Moller, S. Osher, G. Sapiro, and J. Xin, *A convex model for nonnegative matrix factorization and dimensionality reduction on physical space*, IEEE Trans. Image Process., 21 (2012), pp. 3239–3252, https://doi.org/10.1109/TIP.2012.2190081.

[12] N. Gillis, *Robustness analysis of Hottopixx, a linear programming model for factoring nonnegative matrices*, SIAM J. Matrix Anal. Appl., 34 (2013), pp. 1189–1212, https://doi.org/10.1137/120900629.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

[13] N. Gillis, *Successive nonnegative projection algorithm for robust nonnegative blind source separation*, SIAM J. Imaging Sci., 7 (2014), pp. 1420–1450, https://doi.org/10.1137/130946782.

[14] N. Gillis, *Successive projection algorithm robust to outliers*, in Proceedings of the IEEE International Workshop on Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP 2019), IEEE, 2019, pp. 331–335.

[15] N. Gillis, *Nonnegative Matrix Factorization*, SIAM, Philadelphia, 2020.

[16] N. Gillis, D. Kuang, and H. Park, *Hierarchical clustering of hyperspectral images using rank-two nonnegative matrix factorization*, IEEE Trans. Geosci. Remote Sens., 53 (2014), pp. 2066–2078, https://doi.org/10.1109/TGRS.2014.2352857.

[17] N. Gillis and R. Luce, *Robust near-separable nonnegative matrix factorization using linear optimization*, J. Mach. Learn. Res., 15 (2014), pp. 1249–1280.

[18] N. Gillis and W.-K. Ma, *Enhancing pure-pixel identification performance via preconditioning*, SIAM J. Imaging Sci., 8 (2015), pp. 1161–1186, https://doi.org/10.1137/140994915.

[19] N. Gillis and S. A. Vavasis, *Fast and robust recursive algorithms for separable nonnegative matrix factorization*, IEEE Trans. Pattern Anal. Mach. Intell., 36 (2014), pp. 698–714, https://doi.org/10.1109/TPAMI.2013.226.

[20] N. Gillis and S. A. Vavasis, *Semidefinite programming based preconditioning for more robust near-separable nonnegative matrix factorization*, SIAM J. Optim., 25 (2015), pp. 677–698, https://doi.org/10.1137/130940670.

[21] J. Jin, Z. T. Ke, and S. Luo, *Mixed membership estimation for social networks*, J. Econometrics, 239 (2024), 105369, https://doi.org/10.1016/j.jeconom.2022.12.003.

[22] J. Jin, Z. T. Ke, G. Moryoussef, J. Tang, and J. Wang, *Improved algorithm and bounds for successive projection*, in Proceedings of the International Conference on Learning Representations (ICLR), 2024.

[23] D. Kuang and H. Park, *Fast rank-2 nonnegative matrix factorization for hierarchical document clustering*, in Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, 2013, pp. 739–747.

[24] D. D. Lee and H. S. Seung, *Learning the parts of objects by non-negative matrix factorization*, Nature, 401 (1999), pp. 788–791, https://doi.org/10.1038/44565.

[25] Y. Li, D. M. Sima, S. V. Cauter, A. R. Croitor Sava, U. Himmelreich, Y. Pi, and S. Van Huffel, *Hierarchical non-negative matrix factorization (hNMF): A tissue pattern differentiation method for glioblastoma multiforme diagnosis using MRSI*, NMR Biomed., 26 (2013), pp. 307–319, https://doi.org/10.1002/nbm.2850.

[26] W.-K. Ma, J. M. Bioucas-Dias, T.-H. Chan, N. Gillis, P. Gader, A. J. Plaza, A. Ambikapathi, and C.-Y. Chi, *A signal processing perspective on hyperspectral unmixing: Insights from remote sensing*, IEEE Signal Process. Mag., 31 (2014), pp. 67–81, https://doi.org/10.1109/MSP.2013.2279731.

[27] X. Mao, P. Sarkar, and D. Chakrabarti, *Estimating mixed memberships with sharp eigenvector deviations*, J. Amer. Statist. Assoc., 116 (2021), pp. 1928–1940, https://doi.org/10.1080/01621459.2020.1751645.

[28] T. Mizutani, *Ellipsoidal rounding for nonnegative matrix factorization under noisy separability*, J. Mach. Learn. Res., 15 (2014), pp. 1011–1039.

[29] T. Mizutani, *Refinement of hottopixx method for nonnegative matrix factorization under noisy separability*, SIAM J. Matrix Anal. Appl., 43 (2022), pp. 1029–1057, https://doi.org/10.1137/21M1442206.

[30] N. Nadisic, N. Gillis, and C. Kervazo, *Smoothed separable nonnegative matrix factorization*, Linear Algebra Appl., 676 (2023), pp. 174–204, https://doi.org/10.1016/j.laa.2023.07.013.

[31] J. M. Nascimento and J. M. Dias, *Vertex component analysis: A fast algorithm to unmix hyperspectral data*, IEEE Trans. Geosci. Remote Sens., 43 (2005), pp. 898–910, https://doi.org/10.1109/TGRS.2005.844293.

[32] K. Neymeyr and M. Sawall, *On the set of solutions of the nonnegative matrix factorization problem*, SIAM J. Matrix Anal. Appl., 39 (2018), pp. 1049–1069, https://doi.org/10.1137/17M1118439.

[33] C. Pasquini, *Near infrared spectroscopy: A mature analytical technique with new perspectives–A review*, Anal. Chim. Acta, 1026 (2018), pp. 8–36, https://doi.org/10.1016/j.aca.2018.04.004.

[34] B. Recht, C. Re, J. Tropp, and V. Bittorf, *Factoring nonnegative matrices with linear programs*, in Proceedings of Advances in Neural Information Processing Systems (NIPS), 2012, pp. 1214–1222.

[35] N. Sauwen, D. M. Sima, S. Van Cauter, J. Veraart, A. Leemans, F. Maes, U. Himmelreich, and S. Van Huffel, *Hierarchical non-negative matrix factorization to charac-*

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.

Downloaded 09/23/25 to 193.190.208.38 by Nicolas Gillis (nicolas.gillis@umons.ac.be). Redistribution subject to SIAM license or copyright; see https://epubs.siam.org/terms-privacy

terize brain tumor heterogeneity using multi-parametric MRI, NMR Biomed., 28 (2015), pp. 1599–1624, https://doi.org/10.1002/nbm.3413.

[36] T. Virtanen, *Monaural sound source separation by nonnegative matrix factorization with temporal continuity and sparseness criteria*, IEEE Trans. Audio Speech Lang. Process., 15 (2007), pp. 1066–1074, https://doi.org/10.1109/TASL.2006.885253.

[37] O. Vu. Thanh, N. Nadisic, and N. Gillis, *Randomized successive projection algorithm*, in Proceedings of the 28ème Colloque sur le traitement du signal et des images (GRETSI), Nancy, 2022, pp. 893–896, https://gretsi.fr/data/colloque/pdf/2022_nadisic1025.pdf.

[38] R. Wu, W.-K. Ma, Y. Li, A. M.-C. So, and N. D. Sidiropoulos, *Probabilistic simplex component analysis*, IEEE Trans. Signal Process., 70 (2021), pp. 582–599, https://doi.org/10.1109/TSP.2021.3133690.

Copyright © by SIAM. Unauthorized reproduction of this article is prohibited.