Matrix rigidity and the ill-posedness of Robust PCA and matrix completion*

Jared Tanner[†‡] Andrew Thompson[§] Simon Vary[†]

November 14, 2018

Abstract

Robust Principal Component Analysis (PCA) (Candès et al., 2011) and low-rank matrix completion (Recht et al., 2010) are extensions of PCA to allow for outliers and missing entries respectively. It is well-known that solving these problems requires a low coherence between the low-rank matrix and the canonical basis, since in the extreme cases – when the low-rank matrix we wish to recover is also sparse – there is an inherent ambiguity. However, the well-posedness issue in both problems is an even more fundamental one: in some cases, both Robust PCA and matrix completion can fail to have any solutions at due to the set of low-rank plus sparse matrices not being closed, which in turn is equivalent to the notion of the matrix rigidity function not being lower semicontinuous (Kumar et al., 2014). By constructing infinite families of matrices, we derive bounds on the rank and sparsity such that the set of low-rank plus sparse matrices is not closed. We also demonstrate numerically that a wide range of non-convex algorithms for both Robust PCA and matrix completion have diverging components when applied to our constructed matrices. An analogy can be drawn to the case of sets of higher order tensors not being closed under canonical polyadic (CP) tensor rank, rendering the best low-rank tensor approximation unsolvable (Silva and Lim, 2008) and hence encourage the use of multilinear tensor rank (De Lathauwer, 2000).

1 Introduction

Principal Component Analysis (PCA) plays a crucial role in the analysis of high-dimensional data [43, 38, 1, 18] and is a widely used dimensionality reduction technique [23, 26, 36, 33]. It involves solving a low-rank approximation which can be easily computed for moderate size problems [13] by computing the singular value decomposition (SVD), or for larger problem sizes using notions of sketching to compute leading portions of the SVD [22, 14, 47]. Over the last decade PCA has been extended to allow for missing data (matrix completion) or data with either corrupted or few entries inconsistent with a low-rank model (Robust PCA). In this manuscript we show that the set of matrices which are the sum of low-rank and sparse matrices is not closed for a range of rank, sparsity, and matrix dimensions; see Theorem 1.1. Consequently there are a number of algorithms which seek such a decomposition where the constituents diverge while at the same time the sum of the matrices converges, see Section 3. We

*This publication is based on work partially supported by: the EPSRC I-CASE studentship (voucher 15220165) in partnership with Leonardo, and The Alan Turing Institute through EPSRC (EP/N510129/1)
[†]Mathematical Institute, University of Oxford, Woodstock Road, Oxford OX2 6GG, UK
[‡]The Alan Turing Institute, British Library, London NW1 2DB, UK
[§]National Physics Laboratory, Hampton Road, Teddington TW11 0LW, UK

thereby highlight a previously unknown issue practitioners might experience using these techniques. The situation is analogous to the lack of closedness for Tensor CP decomposition rank [25, 24] which motivates the notions of multilinear rank approximation [11].

1.1 Prior work

Robust PCA (RPCA) solves a low-rank plus sparse matrix approximation with the sparse component allowing for few but arbitrarily large corruptions in the low-rank structure; that is, a matrix $M \in \mathbb{R}^{m\times n}$ is decomposed into a low-rank matrix L plus a sparse matrix S

$$\min_{X\in\mathbb{R}^{m\times n}} \|X - M\|_F, \quad \text{s.t.} \quad X \in \mathrm{LS}_{m,n}(r, s), \tag{1.1}$$

where $\mathrm{LS}_{m,n}(r, s)$ is the set of $m \times n$ matrices that can be expressed as a rank r matrix L plus a sparsity s matrix S

$$\mathrm{LS}_{m,n}(r, s) = \left\{L + S \in \mathbb{R}^{m\times n} : \operatorname{rank}(L) \le r,\ \|S\|_0 \le s\right\}.$$

We omit the subscript and write $\mathrm{LS}(r, s)$ where the matrix size is implied from the context and use only a single subindex $\mathrm{LS}_n(r, s)$ to denote sets of square matrices $\mathrm{LS}_{n,n}(r, s)$. Allowing the addition of a sparse matrix to the low-rank matrix can be viewed as modelling globally correlated structure in the low-rank component while allowing local inconsistencies, innovations, or corruptions. Exemplar applications of this model include image restoration [20], hyperspectral image denoising [17, 10, 45], face detection [32, 48], acceleration of dynamic MRI data acquisition [35, 49], analysis of medical imagery [2, 16], separation of moving objects in at otherwise static scene [4], and target detection [34, 39] .

Solving Robust PCA as formulated in (1.1) is an NP-hard problem in general. Provable solutions for the problem were first provided in [6, 9] by solving the convex relaxation of the problem

$$\min_{L\in\mathbb{R}^{m\times n}} \|L\|_* + \lambda\|S\|_1, \quad \text{s.t.} \quad M = L + S, \tag{1.2}$$

where $\|\cdot\|_*$ denotes the Schatten 1-norm[1] of a matrix (sum of its singular values) and $\|\cdot\|_1$ denotes the l_1 norm of a vectorised matrix (sum of absolute values of its entries). In [6], authors show that exact decomposition of a low-rank plus sparse matrix is possible for randomly chosen sparsity locations even for the case of the sparsity level s being a fixed fraction αmn with $\alpha \in (0, 1)$. The work of [9] takes a deterministic approach in which corrupted entries can have arbitrary locations but must be sufficiently spread such that the sparsity fraction of each row and column does not exceed α. In both the works of [6] and [9], as well as subsequent extensions, it is common to impose conditions on the singular vectors of the low-rank component being sufficiently uncorrelated with the canonical basis.

Robust PCA is closely related to the problem of recovering a low-rank matrix from incomplete observations referred to as matrix completion [37]. The main difference between the two is that, in the case of a matrix completion, the indices of missing entries are known, and the aim is to solve

$$\min_{L\in\mathbb{R}^{m\times n}} \|P_\Omega(L) - P_\Omega(M)\|_F, \quad \text{s.t.} \quad L \in \mathrm{LS}_{m,n}(r, 0),\ |\Omega^c| = s, \tag{1.3}$$

where P_Ω is entry-wise subsampling of observed entries of M with indices in Ω.

Similarly to the case of Robust PCA, matrix completion can be approached by solving a convex relaxation formulation of the problem [7, 8, 37], but there are also a number of algorithms that solve the non-convex formulation directly while also providing recovery guarantees [5, 21, 29, 30, 40, 41, 46]. Such non-convex methods are typically observed to be able to recover matrices with higher ranks than is possible by solving the convex relaxed problem [40].

[1] The Schatten 1-norm is often also referred to as the nuclear norm [37].

1.2 Main contribution

It is well known that the model $\mathrm{LS}_{m,n}(r,s)$ from (1.1) need not have a unique solution without further constraints, such as the singular vectors of the low-rank component being uncorrelated with the canonical basis as quantified by the incoherence condition with parameter μ

$$\max_{i\in\{1,\ldots,r\}} \|U^* e_i\|_2 \leq \sqrt{\frac{\mu r}{m}}, \qquad \max_{i\in\{1,\ldots,r\}} \|V^* e_i\|_2 \leq \sqrt{\frac{\mu r}{n}}, \tag{1.4}$$

where $L = U\Sigma V^*$ is the singular value decomposition of the rank r component L of size $m \times n$. The incoherence condition for small values of μ ensures that left and right singular vectors are well spread out and not sparse [7, 37].

Trivial examples of matrices with non-unique decompositions in $\mathrm{LS}(r,s)$ include any matrix with two nonzero entries in differing rows and columns as they are in $\mathrm{LS}(r,s)$ for any r and s such that $r+s=2$ with the entries of the matrix assigned to the sparse or low-rank components selected arbitrarily. Moreover, completion of a low-rank matrix is impossible for sampling patterns P_Ω that are disjoint from the support of the matrix M, which can be likely for matrices that have few nonzeros. Both of the aforementioned problems are overcome by imposing a low coherence which ensures the singular vectors of the low-rank matrix have most entries being nonzero [9].

Herein we highlight the presence of a more fundamental difficulty: There are matrices for which Robust PCA and matrix completion can have no solution in that their constituents diverge even while the objective is minimized to zero. This is not because of the ambiguity between possible solutions or lack of information about the matrix, but instead because $\mathrm{LS}_{m,n}(r,s)$ is not a closed set. Moreover, this is not an isolated phenomenon, as sequences of $\mathrm{LS}_{m,n}(r,s)$ matrices converging outside of the set can be constructed for a wide range of ranks, sparsities and matrix sizes.

Theorem 1.1 ($\mathrm{LS}_n(r,s)$ is not closed)**.** *The set of low-rank plus sparse matrices $\mathrm{LS}_n(r,s)$ is not closed for $r \geq 1$, $s \geq 1$ provided $(r+1)(s+2) \leq n$, or provided $(r+2)^{3/2}s^{1/2} \leq n$ where s is a multiple of a squared integer.*

Proof. By Theorem 2.1 and Lemma 2.6. □

Theorem 1.1 implies that there are matrices M such that problem (1.1) is ill-posed in that the objective can be decreased to zero with the constituents L and S diverging with unbounded energy. The problem size bounds in Theorem 1.1 allow for matrices with $r = \mathcal{O}(n^l)$ to have number of corruptions of order $s = \mathcal{O}(n^{2-3l})$ for $l \in [0, 1/2]$, which for constant rank allows s to be quadratic in n, and for $l \in (1/2, 1]$ to have the number of corruptions of order $s = \mathcal{O}(n^{(1-l)})$. In Section 1.2.1 we illustrate the non-closedness of $\mathrm{LS}_3(1,1)$ and the consequent ill-posedness of the corresponding Robust PCA and low-rank matrix completion problems.

1.2.1 Simple example of $\mathrm{LS}_3(1,1)$ being open

Consider solving for the optimal $\mathrm{LS}(1,1)$ approximation to the following 3×3 matrix, which is a special case of construction given in [27] in the context of the matrix rigidity function not being lower semicontinuous.

$$\min_{X\in\mathbb{R}^{3\times 3}} \|X - M\|_F, \qquad \text{s.t.} \quad X \in \mathrm{LS}(1,1),$$
$$M = \begin{bmatrix} 0 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix} \tag{1.5}$$

Consider the following sequence of matrices X_ϵ

$$X_\epsilon = \begin{bmatrix} 0 & 1 & 1 \\ 1 & \epsilon & \epsilon \\ 1 & \epsilon & \epsilon \end{bmatrix} \in \text{LS}(1,1)$$
$$= \underbrace{\begin{bmatrix} 1/\epsilon & 1 & 1 \\ 1 & \epsilon & \epsilon \\ 1 & \epsilon & \epsilon \end{bmatrix}}_{L_\epsilon} + \underbrace{\begin{bmatrix} -1/\epsilon & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}}_{S_\epsilon},$$

which can decrease the objective function $\|X_\epsilon - M\|_F = 2\epsilon$ to zero as $\epsilon \to 0$, but at the cost of the constituents L_ϵ and S_ϵ diverging with unbounded energy. Moreover, the sequence which minimizes the error converges to a matrix M lying outside of the feasible set $\text{LS}(1,1)$ and is in the set $\text{LS}(1,2)$ instead. As a consequence, Robust PCA as posed in (1.5) does not have a global minimum. As the objective function is decreased towards zero, the energy of both the low-rank and the sparse components diverge to infinity. Likewise, we could consider an instance of the matrix completion problem (1.3) in which the top left entry of M is missing and a rank 1 approximation is sought. We see that a rank 1 solution cannot be obtained as there does not exist a choice for the top left entry that would reduce the rank of M to 1. However, the sequence L_ϵ decreases the objective arbitrarily close to zero while the energy of the iterates grows without bounds, $\|L_\epsilon\|_F \to \infty$.

1.3 Connection with matrix rigidity

Robust PCA is closely related to the notion of the *matrix rigidity* function which was originally introduced in complexity theory by Valiant [42] and refers to the minimum number of entries of M that must be changed in order to reduce it to rank r or lower.

$$\text{Rig}(M,r) = \min_{S \in \mathbb{R}^{m\times n}} \|S\|_0, \text{ s.t. } \text{rank}(M-S) \le r.^2$$

Matrix rigidity is upper bounded for any $M \in \mathbb{R}^{n\times n}$ and rank r as

$$\text{Rig}(M,r) \le (n-r)^2. \tag{1.6}$$

due to elementary matrix properties [42]. Matrices which achieve this upper bound for every r are referred to as *maximally rigid* and it was only recently showed in [27] how to construct them explicitly, which was a long standing open question originally posed by Valiant in 1977.

Matrix rigidity has important consequences for complexity of linear algebraic circuits but is also of interest for its mathematical properties. The work of [27] also provides an example of the rigidity function not being lower semicontinuous, which implies the set $\text{LS}_{m,n}(1,1)$ is not closed. Here, we generalize the result, providing non-closedness examples for many ranks, sparsities and matrix sizes, and discuss consequences for Robust PCA and matrix completion problems. In Section 2 we prove Theorem 1.1 and in Section 3 we illustrate how this phenomenon can cause several Robust PCA and matrix completion algorithms to diverge.

[2] Note that the original definition [42] works with $\text{rank}(M+S) \le r$. Here, we change the sign to be consistent with RPCA notation, $M = L + S$ and $\text{rank}(L) \le r$.

2 Main result

We extend the example of $\mathrm{LS}_3(1,1)$ with $M_3 \in \mathbb{R}^{3\times 3}$ given in (1.5) by constructing $M_n, N_n \notin \mathrm{LS}_n(r,s)$ and yet for which there exists a sequence of matrices $M_n^{(i)}(\epsilon)$ which are in $\mathrm{LS}_n(r,s)$ and $\lim_{\epsilon\to 0} \|M_n^{(i)} - M_n^{(i)}(\epsilon)\|_F = 0$. Matrix $M_n(\epsilon)$ as in (2.5) demonstrates that $\mathrm{LS}_n(r,s)$ is not closed for $r \le s$ (Lemma 2.2) and matrix $N_n(\epsilon)$ as in (2.11) is constructed for $r > s$ (Lemma 2.3). In both cases we require n to be sufficiently large in terms of r and s.

For the case $r \le s$, consider M_n and $M_n(\epsilon)$ of the following general form

$$M_n = \begin{pmatrix} 0_{r,r} & A \\ B & 0_{n-r,n-r} \end{pmatrix}, \qquad M_n(\epsilon) = \begin{pmatrix} 0_{r,r} & A \\ B & \epsilon B\, A \end{pmatrix}, \tag{2.1}$$

where $A, B^T \in \mathbb{R}^{r\times(n-r)}$ and $0_{k,k}$ denotes the $k \times k$ matrix with all zero entries. These constructed matrices satisfy the following properties.

Lemma 2.1 (General form of M_n)**.** *Let M_n and $M_n(\epsilon)$ be as defined in* (2.1)*. Then $M_n(\epsilon) \in \mathrm{LS}(r,r)$. Furthermore*

$$\lim_{\epsilon\to 0} \|M_n(\epsilon) - M_n\|_F = 0. \tag{2.2}$$

Proof. We can write $M_n(\epsilon)$ in the form

$$\begin{pmatrix} \frac{1}{\epsilon} I_r \\ B \end{pmatrix} \begin{pmatrix} I_r & \epsilon A \end{pmatrix} + \begin{pmatrix} -\frac{1}{\epsilon} I_r & 0 \\ 0 & 0 \end{pmatrix}, \tag{2.3}$$

which shows that $M_n(\epsilon) \in \mathrm{LS}_n(r,r)$. It also follows trivially from the definition (2.1) that (2.2) is satisfied. ☐

Remark 2.1 (Nested property of $\mathrm{LS}(r,s)$ sets)**.** *Note that $\mathrm{LS}(r,s)$ sets form a partially ordered set*

$$\mathrm{LS}(r,s) \subseteq \mathrm{LS}(r',s'), \tag{2.4}$$

for any $r' \ge r$ and $s' \ge s$. As a consequence $M_n(\epsilon) \in \mathrm{LS}_n(r,r)$ implies that also $M_n(\epsilon) \in \mathrm{LS}_n(r,s)$ for $s \ge r$.

With Lemma 2.1 we give the general form of M_n and $M_n(\epsilon)$ such that $M_n(\epsilon) \in \mathrm{LS}_n(r,s)$ for $s \ge r$. It remains to show that, for a more specific choice of A and B, we also have $M_n \notin \mathrm{LS}_n(r,s)$. In particular, we construct M_n and $M_n(\epsilon)$ as follows.

$$\begin{aligned} M_n &= \begin{pmatrix} 0_{r,r} & \beta & A^{(1)} & \dots & A^{(l)} \\ \alpha^T & 0_{k,k} & \dots & \dots & 0_{k,r} \\ B^{(1)} & \vdots & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ B^{(l)} & 0_{r,k} & \dots & \dots & 0_{r,r} \end{pmatrix}, \\ M_n(\epsilon) &= \begin{pmatrix} 0_{r,r} & \beta & A^{(1)} & \dots & A^{(l)} \\ \alpha^T & \epsilon\alpha^T\beta & \dots & \dots & \epsilon\alpha^T A^{(l)} \\ B^{(1)} & \vdots & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ B^{(l)} & \epsilon B^{(l)}\beta & \dots & \dots & \epsilon B^{(l)} A^{(l)} \end{pmatrix}, \end{aligned} \tag{2.5}$$

where $\alpha, \beta \in \mathbb{R}^{r\times k}$ are matrices with all non-zero entries, $A^{(i)}, B^{(i)} \in \mathbb{R}^{r\times r}$ are arbitrary non-singular matrices which may, but need not, be the same, $0_{a,b}$ and $\mathbb{1}_{a,b}$ denote $a \times b$ matrices with all entries equal to zero or one respectively, and we set $l = \lceil (s+1)/2 \rceil$, $k = \lceil l/r \rceil$.

By construction, the matrix size is $n = r(l+1)+k$, due to the l matrices $A^{(i)}$ and $B^{(i)}$ for $i = 1, \dots, l$ each being of size $r \times r$, the top left $r \times r$ zero matrix and k columns of α and β.

Lemma 2.2. $\mathrm{LS}_n(r,s)$ *is not closed for* $1 \le r \le s$ *provided*

$$n \ge r\left(\left\lceil \frac{s+1}{2} \right\rceil + 1\right) + \left\lceil \frac{\lceil (s+1)/2 \rceil}{r} \right\rceil. \tag{2.6}$$

Proof. Take M_n as in (2.5). By Lemma 2.1 there exists a matrix sequence $M_n(\epsilon) \in \mathrm{LS}_n(r,r)$ such that $\|M_n(\epsilon) - M_n\|_F \to 0$ as $\epsilon \to 0$. Since for $r \le s$ we have $\mathrm{LS}_n(r,r) \subseteq \mathrm{LS}_n(r,s)$, it follows also that $M_n(\epsilon) \in \mathrm{LS}_n(r,s)$.

It remains to prove that $M_n \notin \mathrm{LS}_n(r,s)$, which is equivalent to showing $\mathrm{Rig}(M_n, r) > s$. We show that having a sparse component $\|S\|_0 \le s$ is insufficient for $\mathrm{rank}(M_n - S) \le r$, because for any choice of such S with at most s non-zero entries, the matrix $M_n - S$ must have a $(r+1)\times(r+1)$ minor with nonzero determinant implying $\mathrm{rank}(M_n - S) \ge r+1$.

In order to establish $\mathrm{rank}(M_n - S) \ge r+1$ we consider $2l$ minors of M_n each of size $(r+1)\times(r+1)$. For l of these we select minors that include $A^{(i)}$, $i = 1, \dots, l$, along with an additional column from the first r columns and an additional row entry from row index $r+1$ to $k+r$ from M_n; and for the remaining l minors we similarly choose a $B^{(i)}$ and an additional row and column as before.

These minors are of the form C_i as shown in (2.7) where the α_i, β_i are chosen to be different entries from α, β for each $i = 1, \dots, l$. This requires α, β to be of size $r \times k$ for $k = \lceil l/r \rceil$. Recall that, by construction of M_n, the α, β have no zero entries and $A^{(i)}, B^{(i)}$ are each full rank. The C_i are constructed as

$$C_i = \begin{cases} \begin{pmatrix} 0_{r,1} & A^{(i)} \\ \alpha_i & 0_{1,r} \end{pmatrix}, & i = 1, \dots, l, \\ \begin{pmatrix} 0_{1,r} & \beta_{i-l} \\ B^{(i-l)} & 0_{r,1} \end{pmatrix}, & i = l+1, \dots, 2l, \end{cases} \tag{2.7}$$

where $0_{u,v}$ denotes the $u \times v$ matrix with all entries equal to zero. Note that matrices C_i do not have disjoint supports as they have some elements from the top left $r \times r$ submatrix of M_n in common. These are the left r zero entries in the first row of C_i for $i = 1, \dots, l$ and the top r zero entries in the first column of C_i for $i = (l+1, \dots, 2l)$. We refer to these entries as the *intersecting part* of C_i.

We now consider the possible S such that $\mathrm{rank}(M_n - S) = r$ and show that any such S must have at least $2l$ nonzeros, thus $\mathrm{Rig}(M_n, r) \ge 2l$. This follows by noting that although the C_i have intersecting portions, S restricted to the i^{th} subminor associated with C_i will have at least one distinct nonzero per i. Consider the C_i for $i = 1, \dots, l$ associated with α_i and $A^{(i)}$ and let S_i be the corresponding $(r+1)\times(r+1)$ sparsity mask of S. It follows that S_i must have at least one entry in the non-intersecting set otherwise $C_i + S_i$ is of the form

$$C_i + S_i = \begin{vmatrix} | & & & \\ s_i & & A^{(i)} & \\ | & & & \\ \alpha_i & 0 & \dots & 0 \end{vmatrix} = \alpha_i |A^{(i)}| \ne 0, \tag{2.8}$$

which is insufficient for the rank of C_i to become rank deficient; similarly for $i = l+1, \dots, 2l$.

Having shown $\mathrm{Rig}(M_n, r) \geq 2l$ we set $l = \lceil (s+1)/2 \rceil$, which then implies that $M_n \notin \mathrm{LS}_n(r, s)$. By the construction of M_n in this argument we have

$$n \geq r(l+1) + k \tag{2.9}$$

due to the l matrices $A^{(i)}$ and $B^{(i)}$ each of size $r \times r$, the top left $r \times r$ matrix $0_{r,r}$ and k columns of β or rows of α respectively, and by zero padding of the matrix we can arbitrarily increase its size. Substituting $l = \lceil (s+1)/2 \rceil$ and $k = \lceil l/r \rceil$, we conclude that $\mathrm{LS}_n(r, s)$ is not a closed set for $s \geq r \geq 1$ provided

$$n \geq r \left(\left\lceil \frac{s+1}{2} \right\rceil + 1 \right) + \left\lceil \frac{\lceil (s+1)/2 \rceil}{r} \right\rceil . \tag{2.10}$$

□

Turning to the $r > s$ case, we now build upon Lemma 2.3 by constructing matrices N_n and $N_n(\epsilon)$ as

$$N_n = \begin{pmatrix} \hat{M}_{n'} & 0 & \dots & 0 \\ 0 & E^{(1,1)} & \dots & E^{(1,s+1)} \\ \vdots & \vdots & \ddots & \\ 0 & E^{(s+1,1)} & & E^{(s+1,s+1)} \end{pmatrix} = \begin{pmatrix} \hat{M}_{n'} & 0_{n',\,(s+1)(r-s)} \\ 0_{(s+1)(r-s),\,n'} & E \end{pmatrix}, \tag{2.11}$$

$$N_n(\epsilon) = \begin{pmatrix} \hat{M}_{n'}(\epsilon) & 0_{n',\,(s+1)(r-s)} \\ 0_{(s+1)(r-s),\,n'} & E \end{pmatrix}$$

where $E^{(i,j)} \in \mathbb{R}^{(r-s)\times(r-s)}$ are identical full rank matrices and

$$\hat{M}_{n'} = \begin{pmatrix} 0_{s,s} & \beta & A^{(1)} & \dots & A^{(l)} \\ \alpha^T & 0 & \dots & \dots & 0_{1,s} \\ B^{(1)} & \vdots & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ B^{(l)} & 0_{s,1} & \dots & \dots & 0_{s,s} \end{pmatrix}, \quad \hat{M}_{n'}(\epsilon) = \begin{pmatrix} 0_{s,s} & \beta & A^{(1)} & \dots & A^{(l)} \\ \alpha^T & \epsilon\alpha^T\beta & \dots & \dots & \epsilon\alpha^T A^{(l)} \\ B^{(1)} & \vdots & \ddots & & \vdots \\ \vdots & \vdots & & \ddots & \vdots \\ B^{(l)} & \epsilon B^{(l)}\beta & \dots & \dots & \epsilon B^{(l)} A^{(l)} \end{pmatrix}, \tag{2.12}$$

have the same structure as in (2.5) but with r replaced by s and as a result $A^{(i,j)}, B^{(i,j)} \in \mathbb{R}^{s\times s}$, $\alpha, \beta \in \mathbb{R}^s$, $l = \lceil (s+1)/2 \rceil$, so $\hat{M}_{n'} \notin \mathrm{LS}_{n'}(s, s)$ while $\hat{M}_{n'}(\epsilon) \in \mathrm{LS}_{n'}(s, s)$.

By construction, the size of $\hat{M}_{n'}$ is $n' = s(l+1) + 1$ and the size of N_n is $n = n' + (s+1)(r-s)$.

Lemma 2.3. $\mathrm{LS}_n(r, s)$ *is not closed for* $r > s \geq 1$ *provided*

$$n \geq s \left(\left\lceil \frac{s+1}{2} \right\rceil + 1 \right) + 1 + (s+1)(r-s). \tag{2.13}$$

Proof. Consider N_n and $N_n(\epsilon)$ from (2.11). By additivity of rank for block diagonal matrices, $\mathrm{rank}\,(E) = (r-s)$ and $\hat{M}_{n'}(\epsilon) \in \mathrm{LS}_{n'}(s, s)$, we have that $N_n(\epsilon) \in \mathrm{LS}_n(r, s)$.

It remains to show that $N_n \notin \mathrm{LS}_n(r, s)$ by proving that $\mathrm{Rig}(N_n, r) > s$. We show that having a sparse component $\|S\|_0 \leq s$ is insufficient for $\mathrm{rank}(N_n - S) \leq r$, because for any such S, matrix $(N_n - S)$ must have at least one $(r+1) \times (r+1)$ minor with non-zero determinant, implying $\mathrm{rank}(N_n - S) \geq r + 1$.

We consider minors D_i of size $(r+1) \times (r+1)$ by diagonally appending a minor $\hat{C}_i \in \mathbb{R}^{(s+1)\times(s+1)}$ of $\hat{M}_{n'}$ of a similar structure as in (2.7) and the whole i^{th} diagonal block $E^{(i,i)} \in \mathbb{R}^{(r-s)\times(r-s)}$

$$D_i = \begin{pmatrix} \hat{C}_i & 0 \\ 0 & E^{(i,i)} \end{pmatrix}, \qquad i = 1, \dots, s+1. \tag{2.14}$$

Due to matrices $E^{(i,i)}$ being picked from the block diagonal, the intersecting parts of supports between D_i are only the intersecting parts between individual $\hat{C}_i$ as explained in (2.7) in the proof of Lemma 2.2. We will ensure that in order for $\operatorname{rank}(D_i) \leq r$ we require S_i to have at least one non-zero in a part of D_i that is disjoint from D_j for $j \neq i$. Either S_i has at least one non-zero on a zero block or $E^{(i,j)}$ or $\hat{C}_i$. If the non-zero is in a zero block or $E^{(i,j)}$, then these are disjoint which implies at least $s+1$ non-zero entries. On the other hand, if the non-zero is in $\hat{C}^{(i)}$ then at least one entry of E must be changed in the non-intersecting part of $\hat{C}_i$ as argued following equation (2.7). Therefore for every D_i at least one distinct entry per i must be changed using the corresponding sparsity component S_i, and since $i = 1, \ldots, s+1$, we must also change at least $s+1$ entries of N_n. We thus have $\operatorname{Rig}(N_n, r) \geq s+1$.

By the construction of N_n in this argument we have

$$n \geq \underbrace{s(l+1)+1}_{n',\text{ size of } \hat{M}_{n'}} + \underbrace{(s+1)(r-s)}_{\text{size of } \mathbb{1}_{s+1}\otimes N}, \tag{2.15}$$

where the size of $\hat{M}_{n'}$ comes from l times repeating the matrices $A^{(i)}$ and $B^{(i)}$ each of size $s \times s$, the top left $s \times s$ matrix $0_{s,s}$, the β column and α row respectively and $s+1$ times repeating matrix E of size $(r-s)$. By zero padding of the matrix we can arbitrarily increase its size. Substituting $l = \lceil (s+1)/2 \rceil$ gives that $\operatorname{LS}_n(r,s)$ is not a closed set for $r > s$ provided

$$n \geq s\left(\left\lceil \frac{s+1}{2} \right\rceil + 1\right) + 1 + (s+1)(r-s). \tag{2.16}$$

□

The following theorem gives a sufficient lower bound on the matrix size such that both size requirements derived in Lemma 2.2 and Lemma 2.3 are met, thus unifying both results.

Theorem 2.1. *The low-rank plus sparse set* $\operatorname{LS}_n(r,s)$ *is not closed provided* $n \geq (r+1)(s+2)$ *and* $r \geq 1$, $s \geq 1$.

Proof. Suppose $n \geq (r+1)(s+2)$. We show that this is a sufficient condition for the matrix size requirements in (2.6) in Lemma 2.2 and (2.13) in Lemma 2.3 to hold.

We first obtain a sufficient condition on the matrix size in (2.6) in Lemma 2.2, bounding

$$\begin{aligned}
&r\left(\left\lceil \frac{s+1}{2} + 1 \right\rceil\right) + \left\lceil \frac{\lceil (s+1)/2 \rceil}{r} \right\rceil \\
\leq\, &r\left(\frac{s+1}{2} + 2\right) + \left(\frac{1}{r}\right)\left(\frac{s+1}{2} + 1\right) + 1 \\
\leq\, &r\left(\frac{s+5}{2}\right) + \left(\frac{s+5}{2}\right) = (r+1)\left(\frac{s+5}{2}\right) \\
\leq\, &(r+1)(s+2),
\end{aligned} \tag{2.17}$$

where the first inequality in (2.17) comes from an upper bound on the ceiling function $\lceil x \rceil \leq x+1$, the second inequality follows from $r \geq 1$ and the last inequality holds for $s \geq 1$.

We also obtain a sufficient bound condition on the matrix size in (2.13) in Lemma 2.3 of the form

$$\begin{aligned}
&s\left(\left\lceil\frac{s+1}{2}+1\right\rceil\right)+1+(s+1)(r-s)\\
&\le s\left(\frac{s+1}{2}+2\right)+(s+1)(r-s)=-\frac{s^2}{2}+\frac{3}{2}+rs+1\\
&\le (r+1)(s+1)\le (r+1)(s+2).
\end{aligned}\tag{2.18}$$

The first inequality in (2.18) comes from an upper bound on the ceiling function and the second inequality holds for $s \geq 1$.

Combining (2.17), (2.18) with Lemma 2.2 and Lemma 2.3 gives that $\mathrm{LS}_n(r,s)$ is not a closed set for $n \geq (r+1)(s+2)$ and $r \geq 1$, $s \geq 1$. □

2.1 Quadratic sparsity

Note that the condition $n \geq (r+1)(s+1)$ limits the order of r and s; in particular if $r = \mathcal{O}(n^l)$ then $s = \mathcal{O}(n^{1-l})$ which for $l \geq 0$ constrains s to be at most linear in n, $s = \mathcal{O}(n)$. In Lemma 2.4 and Lemma 2.5, we extend the result so that for $r = \mathcal{O}(n^l)$ and $l \leq 1/2$ we obtain $s = \mathcal{O}(n^{2-3l})$ which for constant rank, $l = 0$, allows s to be quadratic $\mathcal{O}(n^2)$.

Lemma 2.4 establishes a lower bound on the rigidity of block matrices in terms of the rigidity of a single block. Lemma 2.5 shows that the sequence $K(\epsilon)$ converging to K is an example of $\mathrm{LS}_n(r, p^2 r)$ not being closed provided $n \geq p\left(r\left(\left\lceil\frac{r+1}{2}\right\rceil+1\right)+1\right)$. Let

$$K=\begin{pmatrix}\hat{M}_{n'}^{(1,1)} & \cdots & \hat{M}_{n'}^{(1,p)}\\ \vdots & \ddots & \vdots\\ \hat{M}_{n'}^{(p,1)} & \cdots & \hat{M}_{n'}^{(p,p)}\end{pmatrix},\quad K(\epsilon)=\begin{pmatrix}\hat{M}_{n'}^{(1,1)}(\epsilon) & \cdots & \hat{M}_{n'}^{(1,p)}(\epsilon)\\ \vdots & \ddots & \vdots\\ \hat{M}_{n'}^{(p,1)}(\epsilon) & \cdots & \hat{M}_{n'}^{(p,p)}(\epsilon)\end{pmatrix}\tag{2.19}$$

where matrices $\hat{M}_{n'}^{(i,j)}(\epsilon) \in \mathrm{LS}_{n'}(r,r)$ and $\hat{M}_{n'}^{(i,j)} \notin \mathrm{LS}_{n'}(r,r)$ are of the same structure as in (2.12) and $\lim_{\epsilon\to 0} K(\epsilon) = K$ where $K \in \mathbb{R}^{(n'p)\times(n'p)}$ is constructed by repeating $\hat{M}_{n'}$ in p row and column blocks.

Lemma 2.4. *For K as in* (2.19)

$$\mathrm{Rig}(K,r) \geq p^2\,\mathrm{Rig}(\hat{M}_{n'},r).\tag{2.20}$$

Proof. Let S be the sparsity matrix corresponding to $\mathrm{Rig}(K,r)$, such that

$$\begin{aligned}
&\mathrm{rank}(K-S)\leq r, \qquad \|S\|_0=\mathrm{Rig}(K,r),\\
&\text{and}\qquad S=\begin{pmatrix}\hat{S}^{(1,1)} & \cdots & \hat{S}^{(1,p)}\\ \vdots & \ddots & \vdots\\ \hat{S}^{(p,1)} & \cdots & \hat{S}^{(p,p)}\end{pmatrix},
\end{aligned}\tag{2.21}$$

where $\hat{S}^{(i,j)} \in \mathbb{R}^{n'\times n'}$ denotes the sparsity matrix used in the place of the $\hat{M}_{n'}^{(i,j)}$ block. A necessary condition for $\mathrm{rank}(K-S) \leq r$ is that also the rank of individual blocks is less than or equal to r, that is

$$\mathrm{rank}(\hat{M}_{n'}-\hat{S}^{(i,j)})\leq r, \qquad \forall i,j\in\{1,\dots,p\}\,.\tag{2.22}$$

By definition of the rigidity function as the minimal sparsity of S such that $\operatorname{rank}(\hat{M}_{n'} - S) \leq r$, we have that

$$\|\hat{S}^{(i,j)}\|_0 \geq \operatorname{Rig}(\hat{M}_{n'}, r). \tag{2.23}$$

Summing over all blocks $i, j \in \{1, \ldots, p\}$ yields the result

$$\|S\|_0 = \sum_{i,j}^{p,p} \|\hat{S}^{(i,j)}\|_0 \geq \sum_{i,j}^{p,p} \operatorname{Rig}(\hat{M}_{n'}, r), \tag{2.24}$$

and consequently that

$$\operatorname{Rig}(K, r) \geq p^2 \operatorname{Rig}(\hat{M}_{n'}, r). \tag{2.25}$$

□

Lemma 2.5. $\mathrm{LS}_n(r, p^2 r)$ *is not closed provided*

$$n \geq p\left(r\left(\left\lceil \frac{r+1}{2} \right\rceil + 1\right) + 1\right)$$

and $r \geq 1$, $p \geq 1$.

Proof. Consider K and $K(\epsilon)$ as in (2.19). Repeating $\hat{M}_{n'} \in \mathrm{LS}_{n'}(r, r)$ p times in row and column blocks does not increase the rank, so $\operatorname{rank}(K(\epsilon)) = r$ and by additivity of sparsity we have that $K(\epsilon) \in \mathrm{LS}_n(r, p^2 r)$. By Lemma 2.4 and $\operatorname{Rig}(\hat{M}_{n'}, r) > r$ we have the strict lower bound on the rigidity of K

$$\operatorname{Rig}(K, r) \geq p^2 \operatorname{Rig}(\hat{M}_{n'}, r) > p^2 r, \tag{2.26}$$

which implies that $K \notin \mathrm{LS}_n(r, p^2 r)$ while $K(\epsilon) \in \mathrm{LS}_n(r, p^2 r)$.

Recall that the size of $\hat{M}_{n'}$ as defined in (2.12) is $n' = r(l+1) + 1$ and, since $\hat{M}_{n'}$ is repeated p times, we obtain

$$n \geq p\left(r(l+1) + 1\right) = p\left(r\left(\left\lceil \frac{r+1}{2} \right\rceil + 1\right) + 1\right), \tag{2.27}$$

where the inequality comes from zero padding of the matrix to arbitrarily expand its size. □

Lemma 2.6. *The low-rank plus sparse set* $\mathrm{LS}_n(r, s)$ *is not closed provided*

$$n \geq (r+2)^{3/2} s^{1/2}$$

and $r \geq 1$, $p \geq 1$.

Proof. We weaken the condition of Lemma 2.5 and show that it suffices to have $n \geq (r+2)^{3/2} s^{1/2}$ for $\mathrm{LS}_n(r, s)$ not closed by substituting $s = p^2 r$

$$p\left(r\left(\left\lceil \frac{r+1}{2} \right\rceil + 1\right) + 1\right) = \left(\frac{s}{r}\right)^{\frac{1}{2}}\left(r\left(\left\lceil \frac{r+1}{2} \right\rceil + 1\right) + 1\right) \tag{2.28}$$

$$\leq s^{1/2}\left(r^{1/2}\left(\frac{r+5}{2}\right) + 1\right) = s^{1/2}\left(\frac{r^{3/2}}{2} + 2r^{1/2} + r^{-1/2}\right) \tag{2.29}$$

$$\leq s^{1/2}\left(\frac{r^{3/2}}{2} + 2r^{1/2} + \frac{3}{2} r^{-1/2}\right) = s^{1/2}\frac{(r+1)(r+2)}{2\sqrt{r}} \tag{2.30}$$

$$\leq s^{1/2}(r+2)^{3/2}, \tag{2.31}$$

where in the first line we substitute $s = p^2 r$, the first inequality comes from an upper bound on the ceiling function, the second inequality follows from $r^{-1/2} \leq \frac{3}{2} r^{-1/2}$, and the last inequality holds for $r \geq 1$. □

2.2 Almost maximally rigid examples of non-closedness

We would like to prove non-closedness of $\mathrm{LS}_n(r,s)$ sets for as high ranks r and sparsities s as possible. There cannot be a maximally rigid sequence converging outside $\mathrm{LS}\left(r,(n-r)^2\right)$ because $\mathrm{LS}\left(r,(n-r)^2\right)$ corresponds to the set of all $\mathbb{R}^{n\times n}$ matrices. Similarly, it is necessary that both $r \geq 1$ and $s \geq 1$ hold since sets of rank r matrices $\mathrm{LS}(r,0)$ and sets of sparsity s matrices $\mathrm{LS}(0,s)$ are both closed. As a consequence, the highest possible rank and sparsity for which we may hope to prove that $\mathrm{LS}_n(r,s)$ is not closed corresponds to one strictly less than the maximal rigidity bound, i.e. $\mathrm{LS}\left(r,(n-r)^2-1\right)$ for $r \geq 1$ and also $s=(n-r)^2-1 \geq 1$.

It is shown in [27] that the matrix rigidity function might not be semicontinuous even for maximally rigid matrices. This translates into the set $\mathrm{LS}_3(1,3)$ not being closed as we have $M(\epsilon) \in \mathrm{LS}_3(1,3)$ which converges to $M \notin \mathrm{LS}_3(1,3)$ by choosing

$$M=\begin{pmatrix} a & b & c \\ d & e & 0 \\ g & 0 & i \end{pmatrix} \quad \text{and} \quad M(\epsilon)=\begin{pmatrix} a & b & c \\ d & e & \epsilon cd \\ g & \epsilon bg & i \end{pmatrix}. \tag{2.32}$$

It is easy to check that for a general choice of $\{a,\ldots,i\}$, M is maximally rigid with $\mathrm{Rig}(M,1)=4$. However, $\mathrm{Rig}\left(M(\epsilon),1\right)=3$ since $M(\epsilon)$ can be expressed in the following way

$$M(\epsilon)=\begin{pmatrix} \epsilon^{-1} & b & c \\ d & \epsilon bd & \epsilon cd \\ g & \epsilon bg & \epsilon cg \end{pmatrix}+\begin{pmatrix} a-\epsilon^{-1} & 0 & 0 \\ 0 & e-\epsilon bd & 0 \\ 0 & 0 & i-\epsilon cg \end{pmatrix}. \tag{2.33}$$

We therefore have that $\mathrm{LS}_3(1,3)$ is not a closed set, which is the optimal result with the highest possible sparsity for sets of rank 1 matrices of size 3×3. We pose the question as to whether this result can be generalized and the following conjecture holds.

Conjecture 2.1 (Almost maximally rigid non-closedness)**.** *The low-rank plus sparse set* $\mathrm{LS}_n(r,s)$ *is not closed provided*

$$n \geq r+(s+1)^{1/2}, \tag{2.34}$$

for $s\in[1,(n-1)^2-1]$ *and* $r\in[1,n-2]$.

3 Numerical examples with divergent Robust PCA and matrix completion

Theorem 1.1 and the constructions in Section 2 indicate that there are matrices for which Robust PCA and matrix completion, as stated in (1.1) and (1.3) respectively, are not well defined. In particular, the objective can be driven to zero while the components diverge with unbounded norms. Herein we give examples of two simple matrices which are of a similar construction to M in (1.5),

$$M^{(1)}=\begin{pmatrix} 2 & -1 & -1 \\ -1 & 0 & 0 \\ -1 & 0 & 0 \end{pmatrix}, \quad M^{(2)}=\begin{pmatrix} 1 & -2 & -2 \\ -2 & 0 & 0 \\ -2 & 0 & 0 \end{pmatrix},$$

which are not in $\mathrm{LS}(1,1)$, but can be approximated by an arbitrarily close $M_\epsilon^{(1)}, M_\epsilon^{(2)} \in \mathrm{LS}(1,1)$, and for which popular RPCA and MC algorithms exhibit this divergence. This is analogous to the problem of

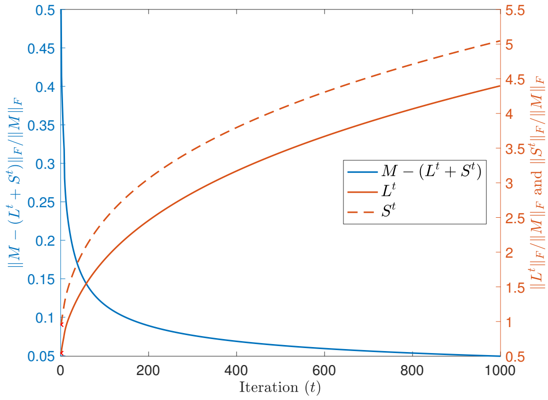

(a) FastGD [50] applied to $M^{(1)}$.

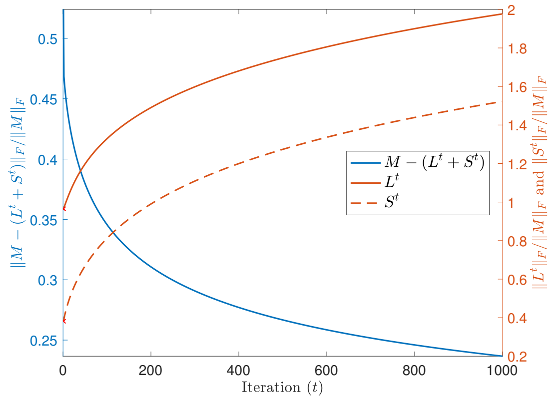

(b) AltMin [19] applied to $M^{(2)}$.

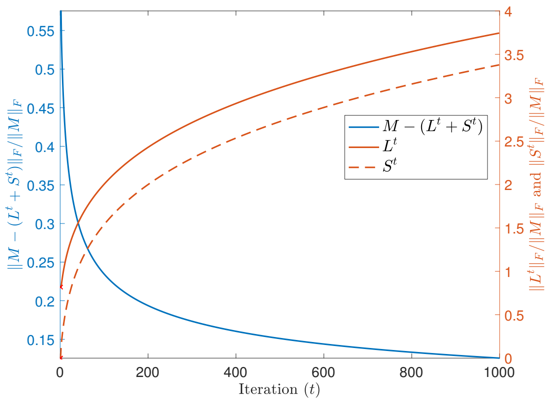

(c) AltProj [15] applied to $M^{(2)}$.

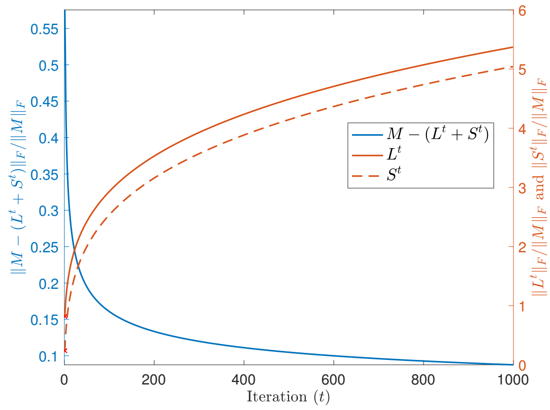

(d) GoDec [31] applied to $M^{(2)}$.

Figure 1: Solving for an LS$(1,1)$ approximation to $M^{(1)}$ and $M^{(2)}$ using four non-convex Robust PCA algorithms. Despite the norm of the residual $\|M-(L^t+S^t)\|_F$ converging to zero, norms of the constituents L^t, S^t diverge. We set algorithms parameters $r=1, s=1$ where possible. For FastGD we set $\lambda=3.23$ and stepsize $\eta=1/6$ which corresponds to choosing $s=1$. For GoDec we set the low-rank projection precision parameter to be 10.

diverging components for CP-rank decomposition of higher order tensors which is especially pronounced for algorithms employing alternating search between individual components [12].

Non-convex algorithms for solving the Robust PCA problem (1.1) are typically observed to be faster than convex relaxations of the problem and often are able to recover matrices with higher ranks than possible by solving the convex relaxation (1.2). We explore the performance of four widely considered non-convex Robust PCA algorithms: Fast Robust PCA via Gradient Descent (FastGD) [50], Alternating Minimization (AltMin) [19], Alternating Projection (AltProj) [15], and Go Decomposition (GoDec) [31] applied to $M^{(1)}$ or $M^{(2)}$ with algorithm parameters set to rank $r=1$ and sparsity $s=1$. In each case Figure 1 shows the convergence of the residual $\min_{X\in\mathbb{R}^{m\times n}}\|X-M\|_F$ to zero while the norms of the constituents of $M=L+S$ diverge.

Convex relaxations of RPCA such as posed in (1.2) do not suffer from the divergence of constituents as shown in Figure 1 due to their explicit minimization of their norms. However, they suffer from sub-

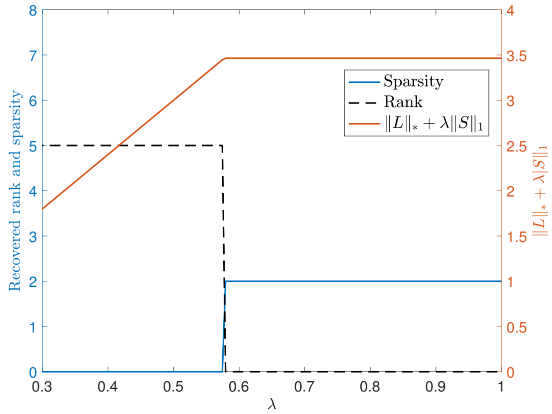

(a) PCP [6] applied to $M^{(1)}$.

(b) IALM [28] applied to $M^{(1)}$.

Figure 2: Recovered ranks and sparsities using two convex Robust PCA algorithms applied to $M^{(1)}$ with varying choice of λ. Both PCP and IALM do not recover the $r = 1, s = 1$ solution for any λ. IALM recovers solutions with overspecified degrees of freedom $r = 2, s = 5$ for λ roughly $1/2$.

optimal performance. Figure 2 depicts recovered ranks, sparsities and their convex relaxations based on choice for λ of $M^{(1)}$ for Principal Component Pursuit by Alternating Directions (PCP) [6] and Inexact Augmented Lagrangian Method (IALM) [28]. For both PCP and IALM, as the regularization parameter λ is increased from near zero it first produced a solution with $r = 0$ and $s = 5$, then at approximately $\lambda = 1/2$ transitions to solutions with overspecified degrees of freedom $r = 2$ and $s = 5$, and then for large values of λ gives solutions with $r = 2$ and $s = 0$. It is interesting to note that for these convex relaxations of RPCA there were no values of λ that would produce a solutions with $r = 1$ and $s = 1$ which are the parameters for which the non-convex RPCA algorithms diverge. In contrast, the aforementioned non-convex algorithms for RPCA applied to $M^{(1)}$ converge to zero residual with bounded constituents for the rank and sparsity parameters generated by PCP and IALM.

Similar to the divergence of the non-convex RPCA algorithms, non-convex matrix completion algorithms applied to $M^{(1)}$ with only the top left, index $(1, 1)$, entry missing can diverge[3]. Figure 3 depicts the residual error converging to zero and energy of the recovered low-rank matrix diverging for four exemplar non-convex algorithms: Power Factorization (PF) [21], Low-Rank Matrix Fitting (LMaFit) [46], Conjugate Gradient Iterative Hard Thresholding (CGIHT) [3] and Alternating Steepest Descent (ASD) [41].

4 Conclusion

This work brings to attention an overlooked issue in Robust PCA and matrix completion: that both problems can be ill-posed because the set of low-rank plus sparse matrices is not closed without further conditions being set on the constituent matrices. It remains to be determined what fraction of the set $L_{m,n}(r, s)$ is open, or similarly what fraction has constituents whose norm exceeds a prescribed threshold to ensure well conditioning; it should be noted that in the case of Tensor CP rank the fraction of the space of tensors with unbounded constituent energy is a positive measure [12]. It also remains to determine what is the maximal matrix size n, as a function of (r, s), such that the set $\mathrm{LS}_n(r, s)$ is open. We give

[3]It is required to provide the algorithm with an initial guess that does not have 0 as the top left entry.

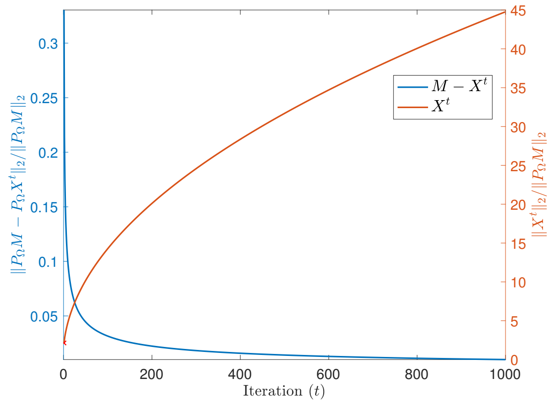

(a) PF [21] applied to $M^{(1)}$.

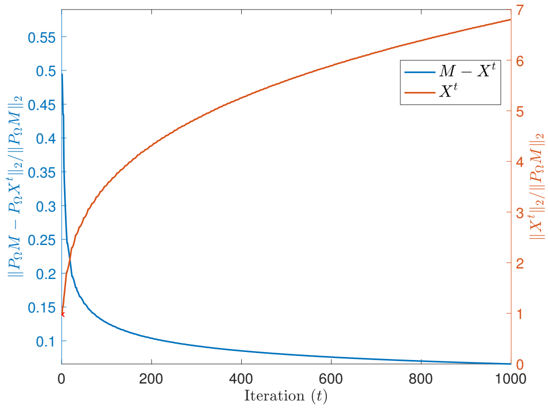

(b) LMaFit [46] applied to $M^{(1)}$.

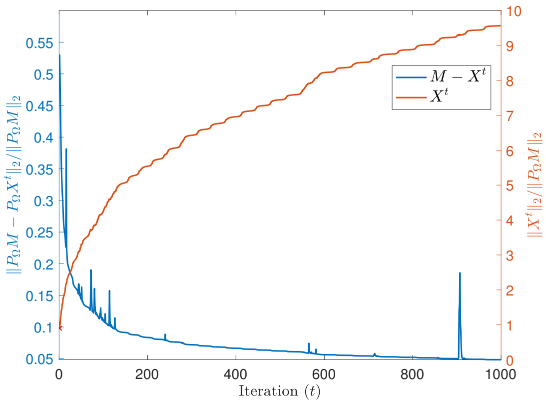

(c) CGIHT with restarts [3] applied to $M^{(1)}$.

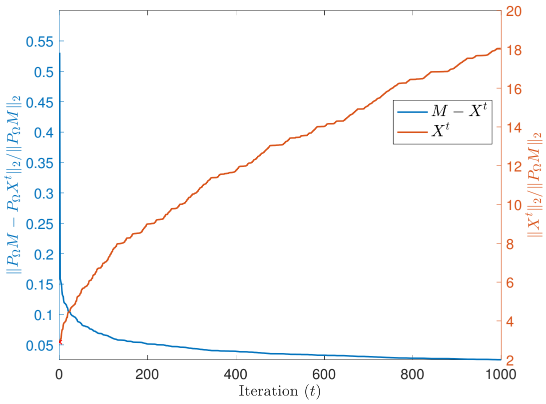

(d) ASD [41] applied to $M^{(1)}$.

Figure 3: Recovery of $M^{(1)}$ given a rank 1 constraint by four non-convex matrix completion algorithm. Despite the norm of the residual $\|y - P_\Omega(X^t)\|_F$ converging to zero, the norm of the recovered matrix X^t diverges.

lower bound of $n(r,s) \geq (r+1)(s+2)$ and $n(r,s) \geq (r+2)^{(3/2)}s^{1/2}$ in Theorem 1.1 and conjecture the best attainable bound is achieved at $n(r,s) \geq r + (s+1)^{1/2}$ using bounds on maximum matrix regidity, see Conjecture 2.1. Moreover, we note that there are references in the literature [19, 44] which reference the use of a restricted isometry property for $\mathrm{LS}_n(r,s)$ in order to prove recovery of RPCA using non-convex algorithms. A consequence of our result is that the lower RIP bound is not satisfied for some $M \in \mathrm{LS}(r,s)$ unless further restrictions are imposed on the constituents such as bounds on the energy of L and S which compose M.

References

[1] H. Abdi and L. J. Williams, *Principal component analysis*, Wiley Interdisciplinary Reviews: Computational Statistics, 2 (2010), pp. 433–459.

[2] S. H. Baete, J. Chen, Y.-C. Lin, X. Wang, R. Otazo, and F. E. Boada, *Low Rank plus Sparse*

decomposition of ODFs for improved detection of group-level differences and variable correlations in white matter, NeuroImage, 174 (2018), pp. 138–152.

[3] J. D. Blanchard, J. Tanner, and K. Wei, *CGIHT: conjugate gradient iterative hard thresholding for compressed sensing and matrix completion*, Information and Inference, (2015), p. iav011.

[4] T. Bouwmans, A. Sobral, S. Javed, S. K. Jung, and E.-H. Zahzah, *Decomposition into low-rank plus additive matrices for background/foreground separation: A review for a comparative evaluation with a large-scale dataset*, Computer Science Review, 23 (2017), pp. 1–71.

[5] J.-F. Cai, E. J. Candès, and Z. Shen, *A Singular Value Thresholding Algorithm for Matrix Completion*, SIAM Journal on Optimization, 20 (2010), pp. 1956–1982.

[6] E. J. Candès, X. Li, Y. Ma, and J. Wright, *Robust principal component analysis?*, Journal of the ACM, 58 (2011), pp. 1–37.

[7] E. J. Candès and B. Recht, *Exact Matrix Completion via Convex Optimization*, Foundations of Computational Mathematics, 9 (2009), pp. 717–772.

[8] E. J. Candes and T. Tao, *The Power of Convex Relaxation: Near-Optimal Matrix Completion*, IEEE Transactions on Information Theory, 56 (2010), pp. 2053–2080.

[9] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky, *Rank-Sparsity Incoherence for Matrix Decomposition*, SIAM Journal on Optimization, 21 (2011), pp. 572–596.

[10] Y. Chen, Y. Guo, Y. Wang, D. Wang, C. Peng, and G. He, *Denoising of Hyperspectral Images Using Nonconvex Low Rank Matrix Approximation*, IEEE Transactions on Geoscience and Remote Sensing, 55 (2017), pp. 5366–5380.

[11] L. De Lathauwer, B. De Moor, and J. Vandewalle, *On the Best Rank-1 and Rank-(R 1 , R 2 ,. . ., R N) Approximation of Higher-Order Tensors*, SIAM Journal on Matrix Analysis and Applications, 21 (2000), pp. 1324–1342.

[12] V. de Silva and L.-h. Lim, *Tensor Rank and the Ill-Posedness of the Best Low-Rank Approximation Problem*, SIAM Journal on Matrix Analysis and Applications, 30 (2008), pp. 1084–1127.

[13] J. W. Demmel, *Applied Numerical Linear Algebra*, Society for Industrial and Applied Mathematics, jan 1997.

[14] P. Drineas, R. Kannan, and M. W. Mahoney, *Fast Monte Carlo Algorithms for Matrices II: Computing a Low-Rank Approximation to a Matrix*, SIAM Journal on Computing, 36 (2006), pp. 158–183.

[15] A. Dutta, F. Hanzely, and P. Richtárik, *A Nonconvex Projection Method for Robust PCA*, (2018), pp. 1–24.

[16] H. Gao, J.-F. Cai, Z. Shen, and H. Zhao, *Robust principal component analysis-based four-dimensional computed tomography*, Physics in Medicine and Biology, 56 (2011), pp. 3181–3198.

[17] A. Gogna, A. Shukla, H. K. Agarwal, and A. Majumdar, *Split Bregman algorithms for sparse / joint-sparse and low-rank signal recovery: Application in compressive hyperspectral imaging*, in 2014 IEEE International Conference on Image Processing (ICIP), IEEE, oct 2014, pp. 1302–1306.

[18] C. Goodall and I. T. Jolliffe, *Principal Component Analysis*, Technometrics, 30 (1988), p. 351.

[19] Q. Gu and Z. Wang, *Low-Rank and Sparse Structure Pursuit via Alternating Minimization*, International Conference on Artificial Intelligence and Statistics, 51 (2016).

[20] S. Gu, L. Zhang, W. Zuo, and X. Feng, *Weighted Nuclear Norm Minimization with Application to Image Denoising*, in 2014 IEEE Conference on Computer Vision and Pattern Recognition, no. 2, IEEE, jun 2014, pp. 2862–2869.

[21] J. Haldar and D. Hernando, *Rank-Constrained Solutions to Linear Matrix Equations Using PowerFactorization*, IEEE Signal Processing Letters, 16 (2009), pp. 584–587.

[22] N. Halko, P. G. Martinsson, and J. A. Tropp, *Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions*, SIAM Review, 53 (2011), pp. 217–288.

[23] E. Han, P. Carbonetto, R. E. Curtis, Y. Wang, J. M. Granka, J. Byrnes, K. Noto, A. R. Kermany, N. M. Myres, M. J. Barber, K. A. Rand, S. Song, T. Roman, E. Battat, E. Elyashiv, H. Guturu, E. L. Hong, K. G. Chahine, and C. A. Ball, *Clustering of 770,000 genomes reveals post-colonial population structure of North America*, Nature Communications, 8 (2017), p. 14238.

[24] F. L. Hitchcock, *The Expression of a Tensor or a Polyadic as a Sum of Products*, Journal of Mathematics and Physics, 6 (1927), pp. 164–189.

[25] F. L. Hitchcock, *Multiple Invariants and Generalized Rank of a P-Way Matrix or Tensor*, Journal of Mathematics and Physics, 7 (1928), pp. 39–79.

[26] L.-C. Hsu, C.-Y. Huang, Y.-H. Chuang, H.-W. Chen, Y.-T. Chan, H. Y. Teah, T.-Y. Chen, C.-F. Chang, Y.-T. Liu, and Y.-M. Tzou, *Accumulation of heavy metals and trace elements in fluvial sediments received effluents from traditional and semiconductor industries*, Scientific Reports, 6 (2016), p. 34250.

[27] A. Kumar, S. V. Lokam, V. M. Patankar, and M. N. J. Sarma, *Using Elimination Theory to Construct Rigid Matrices*, computational complexity, 23 (2014), pp. 531–563.

[28] O. Kuybeda, G. A. Frank, A. Bartesaghi, M. Borgnia, S. Subramaniam, and G. Sapiro, *A collaborative framework for 3D alignment and classification of heterogeneous subvolumes in cryo-electron tomography*, Journal of Structural Biology, 181 (2013), pp. 116–127.

[29] A. Kyrillidis and V. Cevher, *Matrix Recipes for Hard Thresholding Methods*, Journal of Mathematical Imaging and Vision, 48 (2014), pp. 235–265.

[30] K. Lee and Y. Bresler, *ADMiRA: Atomic Decomposition for Minimum Rank Approximation*, IEEE Transactions on Information Theory, 56 (2010), pp. 4402–4416.

[31] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, *Robust Recovery of Subspace Structures by Low-Rank Representation*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 171–184.

[32] X. Luan, B. Fang, L. Liu, W. Yang, and J. Qian, *Extracting sparse error of robust PCA for face recognition in the presence of varying illumination and occlusion*, Pattern Recognition, 47 (2014), pp. 495–508.

[33] J. Miehlbradt, A. Cherpillod, S. Mintchev, M. Coscia, F. Artoni, D. Floreano, and S. Micera, *Data-driven bodymachine interface for the accurate control of drones*, Proceedings of the National Academy of Sciences, 115 (2018), pp. 7913–7918.

[34] O. Oreifej, X. Li, and M. Shah, *Simultaneous Video Stabilization and Moving Object Detection in Turbulence*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 35 (2013), pp. 450–462.

[35] R. Otazo, E. Candès, and D. K. Sodickson, *Low-rank plus sparse matrix decomposition for accelerated dynamic MRI with separation of background and dynamic components*, Magnetic Resonance in Medicine, 73 (2015), pp. 1125–1136.

[36] C. Plesa, A. M. Sidore, N. B. Lubock, D. Zhang, and S. Kosuri, *Multiplexed gene synthesis in emulsions for exploring protein functional landscapes*, Science, 359 (2018), pp. 343–347.

[37] B. Recht, M. Fazel, and P. a. Parrilo, *Guaranteed Minimum-Rank Solutions of Linear Matrix Equations via Nuclear Norm Minimization*, SIAM Review, 52 (2010), pp. 471–501.

[38] M. Ringnér, *What is principal component analysis?*, Nature Biotechnology, 26 (2008).

[39] D. Sabushimike, S. Na, J. Kim, N. Bui, K. Seo, and G. Kim, *Low-Rank Matrix Recovery Approach for Clutter Rejection in Real-Time IR-UWB Radar-Based Moving Target Detection*, Sensors, 16 (2016), p. 1409.

[40] J. Tanner and K. Wei, *Normalized Iterative Hard Thresholding for Matrix Completion*, SIAM Journal on Scientific Computing, 35 (2013), pp. S104–S125.

[41] ——, *Low rank matrix completion by alternating steepest descent methods*, Applied and Computational Harmonic Analysis, 40 (2016), pp. 417–429.

[42] L. G. Valiant, *Graph-theoretic arguments in low-level complexity*, in Mathematical Foundations of Computer Science 1977, J. Gruska, ed., vol. 53, Springer Berlin Heidelberg, Berlin, Heidelberg, 1977, pp. 162–176.

[43] N. Vaswani, Y. Chi, and T. Bouwmans, *Rethinking PCA for Modern Data Sets: Theory, Algorithms, and Applications [Scanning the Issue]*, Proceedings of the IEEE, 106 (2018), pp. 1274–1276.

[44] A. E. Waters, A. C. Sankaranarayanan, and R. G. Baraniuk, *SpaRCS: Recovering Low-rank and Sparse Matrices from Compressive Measurements*, in Proceedings of the 24th International Conference on Neural Information Processing Systems, no. 2, Granada, Spain, 2011, pp. 1089—-1097.

[45] W. Wei, L. Zhang, Y. Zhang, C. Wang, and C. Tian, *Hyperspectral image denoising from an incomplete observation*, in 2015 International Conference on Orange Technologies (ICOT), IEEE, dec 2015, pp. 177–180.

[46] Z. Wen, W. Yin, and Y. Zhang, *Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm*, Mathematical Programming Computation, 4 (2012), pp. 333–361.

[47] D. P. Woodruff, *Sketching as a Tool for Numerical Linear Algebra*, Foundations and Trends in Theoretical Computer Science, 10 (2014), pp. 1–157.

[48] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Yi Ma, *Robust Face Recognition via Sparse Representation*, IEEE Transactions on Pattern Analysis and Machine Intelligence, 31 (2009), pp. 210–227.

[49] F. Xu, J. Han, Y. Wang, M. Chen, Y. Chen, G. He, and Y. Hu, *Dynamic Magnetic Resonance Imaging via Nonconvex Low-Rank Matrix Approximation*, IEEE Access, 5 (2017), pp. 1958–1966.

[50] X. Yi, D. Park, Y. Chen, and C. Caramanis, *Fast Algorithms for Robust PCA via Gradient Descent*, Advances in Neural Information Processing Systems, (2016).