

Iterative Hard Thresholding for Low-Rank Recovery from Rank-One Projections

Simon Foucart* and Srinivas Subramanian — Texas A&M University

Abstract

A novel algorithm for the recovery of low-rank matrices acquired via compressive linear measurements is proposed and analyzed. The algorithm, a variation on the iterative hard thresholding algorithm for low-rank recovery, is designed to succeed in situations where the standard rank-restricted isometry property fails, e.g. in case of subexponential unstructured measurements or of subgaussian rank-one measurements. The stability and robustness of the algorithm are established based on distinctive matrix-analytic ingredients and its performance is substantiated numerically.

Key words and phrases: Compressive sensing, low-rank recovery, singular value thresholding, rank-one projections, rank-restricted isometry properties.

AMS classification: 15A18, 15A29, 65F10, 94A12.

1 Introduction

This article is concerned with the recovery of matrices $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ of rank $r \ll N := \min\{N_1, N_2\}$ from compressive linear measurements

$$(1) \quad \mathbf{y} = \mathcal{A}(\mathbf{X}) \in \mathbb{R}^m$$

with a number m of measurements scaling like $r \max\{N_1, N_2\}$ instead of $N_1 N_2$. It is known that such a low-rank recovery task can be successfully accomplished when the linear map $\mathcal{A} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^m$ satisfies a rank-restricted isometry property [2], e.g. by performing nuclear norm minimization [11], i.e., by outputting a solution of

$$(2) \quad \underset{\mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}}{\text{minimize}} \|\mathbf{Z}\|_* := \sum_{k=1}^N \sigma_k(\mathbf{Z}) \quad \text{subject to } \mathcal{A}(\mathbf{Z}) = \mathbf{y},$$

or by executing an iterative hard thresholding scheme [8, 10], i.e., by outputting the limit of the sequence $(\mathbf{X}_n)_{n \geq 0}$ defined by $\mathbf{X}_0 = \mathbf{0}$ and

$$(3) \quad \mathbf{X}_{n+1} = H_r(\mathbf{X}_n + \mu \mathcal{A}^*(\mathbf{y} - \mathcal{A}\mathbf{X}_n)).$$

*S. F. partially supported by NSF grants DMS-1622134 and DMS-1664803.

Above, and in the rest of the article,

- $\sigma_1(\mathbf{Z}) \geq \dots \geq \sigma_N(\mathbf{Z})$ denote the singular values of a matrix $\mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}$;
- the linear map $\mathcal{A}^* : \mathbb{R}^m \rightarrow \mathbb{R}^{N_1 \times N_2}$ stands for the adjoint of \mathcal{A} , defined by the property that $\langle \mathcal{A}^*(\mathbf{u}), \mathbf{Z} \rangle_F = \langle \mathbf{u}, \mathcal{A}(\mathbf{Z}) \rangle$ for all $\mathbf{u} \in \mathbb{R}^m$ and all $\mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}$;
- H_r represents the operator of best approximation by matrices of rank at most r , so that $H_r(\mathbf{Z}) = \sum_{\ell=1}^r \sigma_\ell(\mathbf{Z}) \mathbf{u}_\ell \mathbf{v}_\ell^*$ if $\mathbf{Z} = \sum_{\ell=1}^N \sigma_\ell(\mathbf{Z}) \mathbf{u}_\ell \mathbf{v}_\ell^*$ is the singular value decomposition of \mathbf{Z} .

Our focus, however, is put on measurement maps that do not necessarily satisfy the standard rank-restricted isometry property, but rather a modified rank-restricted isometry property featuring the ℓ_1 -norm as an inner norm. Precisely, we consider measurement maps $\mathcal{A} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^m$ with universally bounded rank-restricted isometry ratio

$$(4) \quad \gamma_r := \frac{\beta_r}{\alpha_r} \geq 1,$$

where α_r, β_r denote the optimal constants $\alpha, \beta > 0$ such that

$$(5) \quad \alpha \|\mathbf{Z}\|_F \leq \|\mathcal{A}(\mathbf{Z})\|_1 \leq \beta \|\mathbf{Z}\|_F \quad \text{for all } \mathbf{Z} \in \mathbb{R}^{N_1 \times N_2} \text{ with } \text{rank}(\mathbf{Z}) \leq r.$$

For instance, it was shown in [1]¹ that $\gamma_r \leq 3$ holds with high probability in the case of measurements obtained by rank-one projections of the form²

$$(6) \quad \mathcal{A}(\mathbf{Z})_i = \mathbf{a}_i^* \mathbf{Z} \mathbf{b}_i, \quad i = 1, \dots, m,$$

for some independent standard Gaussian vectors $\mathbf{a}_1, \dots, \mathbf{a}_m \in \mathbb{R}^{N_1}$ and $\mathbf{b}_1, \dots, \mathbf{b}_m \in \mathbb{R}^{N_2}$, provided $m \gtrsim r \max\{N_1, N_2\}$. Moreover, we will briefly justify in Section 5 that, in the case of measurements given by

$$(7) \quad \mathcal{A}(\mathbf{Z})_i = \sum_{k,\ell} \mathcal{A}_{i,(k,\ell)} Z_{k,\ell}, \quad i = 1, \dots, m,$$

for some independent identically distributed mean-zero subexponential random variables $\mathcal{A}_{i,(k,\ell)}$ with variance $1/m$, say, there are constants $c, C > 0$ depending on the subexponential distribution such that $\gamma_r \leq c$ holds with high probability as soon as $m \geq Cr \max\{N_1, N_2\}$.

The article [1] also showed that, under the modified rank-restricted isometry property (5), recovery of rank- r matrices $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ from $\mathbf{y} = \mathcal{A}(\mathbf{X}) \in \mathbb{R}^m$ can be achieved by the nuclear norm minimization (2). Here, we highlight the question:

Under the modified rank-restricted isometry property, is it possible to achieve low-rank recovery by an iterative hard thresholding algorithm akin to (3)?

¹Similar statements also appear in [3] and [4] in the cases $r = 1$ and $r = 2$, respectively.

²These are called rank-one projections because $\mathbf{a}_i^* \mathbf{Z} \mathbf{b}_i$ is the Frobenius inner product $\langle \mathbf{Z}, \mathbf{a}_i \mathbf{b}_i^* \rangle_F = \text{tr}(\mathbf{b}_i \mathbf{a}_i^* \mathbf{Z})$ of the matrix \mathbf{Z} with the rank-one matrix $\mathbf{a}_i \mathbf{b}_i^*$.

We shall answer this question in the affirmative. The argument resembles the one developed in [6] in the context of sparse recovery, but some subtle matrix-analytic refinements are necessary.

The iterative hard thresholding scheme we propose consists in outputting the limit of the sequence $(\mathbf{X}_n)_{n \geq 0}$ defined by $\mathbf{X}_0 = \mathbf{0}$ and

$$(8) \quad \mathbf{X}_{n+1} = H_s(\mathbf{X}_n + \mu_n H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))),$$

with parameters s , t , and μ_n to be given explicitly later. If H_t did not appear, the scheme (8) would intuitively be interpreted as a (sub)gradient descent steps for the function $\mathbf{Z} \mapsto \|\mathbf{y} - \mathcal{A}\mathbf{Z}\|_1$, followed by some singular value thresholding. The appearance of H_t is the main adjustment distinguishing the scheme (8) from the iterative hard thresholding scheme proposed in [6] (it also leads to an altered stepsize μ_n). For the more classical iterative hard thresholding scheme, a similar adjustment was already exploited in [9], where the roles of H_s and H_t were to provide ‘tail’ and ‘head’ approximations, respectively. This is not what justifies the appearance of H_t here.³

Let us now state our most representative result. To reiterate, it provides our newly proposed algorithm with theoretical guarantees for the success of recovery of low-rank matrices acquired via rank-one projections. Other iterative thresholding algorithms do not possess such guarantees, since their analysis exploits the standard rank-restricted isometry property and hence does not apply to rank-one projections. We also emphasize that the result is valid beyond rank-one projections, since it only relies on the modified rank-restricted isometry property (5).

Theorem 1. Let $\mathcal{A} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^m$ be a measurement map satisfying the modified rank-restricted isometry property (5) of order $c_0 r$ with ratio $\gamma_{c_0 r} \leq \gamma$. Then any matrix $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ of rank at most r acquired by $\mathbf{y} = \mathcal{A}(\mathbf{X})$ can be exactly recovered as the limit of the sequence $(\mathbf{X}_n)_{n \geq 0}$ produced by (8) with parameters $s = c_1 r$, $t = c_2 r$, and $\mu_n = \|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1 / \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2$. The constants $c_0, c_1, c_2 > 0$ depend only on γ .

Theorem 1 is a simplified version of our main result (Theorem 5), which covers the more realistic situation where the matrices $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ to be recovered are not exactly low-rank and where the measurements are not perfectly accurate, i.e., $\mathbf{y} = \mathcal{A}\mathbf{X} + \mathbf{e}$ with a nonzero vector $\mathbf{e} \in \mathbb{R}^m$. The full result will be stated and proved in Section 3. In preparation, we collect in Section 2 the matrix-analytic tools that are essential to our argument. In Section 4, we present some modest numerical

³ In fact, the experiments performed in Section 4 suggest that H_t is not necessary, but we were unable to prove the recovery of \mathbf{X} using the algorithm ‘without H_t ’. In the vector case, one avoids H_t by relying on the observation that $H_s(\mathbf{z}) = H_s(\mathbf{z}_T)$ if the index set T contains the support of $H_s(\mathbf{z})$. In the matrix case, one would seemingly require the counterpart statement that $H_s(\mathbf{Z}) = H_s(P_{\mathcal{T}}(\mathbf{Z}))$, $P_{\mathcal{T}}$ denoting the orthogonal projection onto a space \mathcal{T} containing the space $\text{span}\{\mathbf{u}_1 \mathbf{v}_1^*, \dots, \mathbf{u}_s \mathbf{v}_s^*\}$ obtained from the singular value decomposition $\sum_{\ell=1}^N \sigma_{\ell}(\mathbf{Z}) \mathbf{u}_{\ell} \mathbf{v}_{\ell}^*$ of \mathbf{Z} .

Such a statement is unfortunately invalid — take e.g. $\mathbf{Z} = \begin{bmatrix} 3 & 0 \\ 0 & 3c \end{bmatrix}$ with $|c| < 1$, so that $H_1(\mathbf{Z}) = \begin{bmatrix} 3 & 0 \\ 0 & 0 \end{bmatrix}$, and take $\mathcal{T} = \text{span} \left\{ \begin{bmatrix} 1 & 0 \\ 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix} \right\}$, so that $P_{\mathcal{T}}(\mathbf{Z}) = \begin{bmatrix} 3 & c \\ c & c \end{bmatrix}$ and $H_1(P_{\mathcal{T}}(\mathbf{Z})) \neq \begin{bmatrix} 3 & 0 \\ 0 & 0 \end{bmatrix}$ (otherwise $\text{rank} \left(\begin{bmatrix} 0 & c \\ c & c \end{bmatrix} \right) = 1$).

experiments demonstrating some strong points of our iterative hard thresholding algorithm. In Section 5, we close the article with a sketched justification of the modified rank-restricted isometry property for subexponential measurement maps.

2 Matrix-analytic ingredients

The proof of Theorem 1 and of its generalization (Theorem 5) is based on the two technical side results given as Propositions 2 and 3 below. The first side result provides a (loose⁴) quantitative confirmation to the intuitive fact that $\|\mathbf{X} - H_s(\mathbf{Z})\|_F$ approaches $\|\mathbf{X} - \mathbf{Z}\|_F$ as s grows. The (tight) vector version of this assertion, due to [12], was already harnessed by [6] in the context of sparse recovery when the standard restricted isometry property is absent.

Proposition 2. Given matrices $\mathbf{X}, \mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}$ and integers $r \leq s$ with $s + 2r \leq \min\{N_1, N_2\}$, if $\text{rank}(\mathbf{X}) \leq r$, then

$$(9) \quad \|\mathbf{X} - H_s(\mathbf{Z})\|_F \leq \eta\left(\frac{s}{r}\right) \|\mathbf{X} - \mathbf{Z}\|_F, \quad \text{where} \quad \eta(\kappa) := 1 + \sqrt{\frac{8}{\kappa + 1}} \xrightarrow{\kappa \rightarrow \infty} 1.$$

The second side result says that, under the modified rank-restricted isometry property (5), a low-rank matrix \mathbf{X} is well approximated by forming \mathcal{A}^* applied to the sign vector of $\mathcal{A}(\mathbf{X})$ and then truncating its singular value decomposition. A closely connected result can be found in [7, Lemma 3]⁵.

Proposition 3. Given a matrix $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$, integers $r \leq s$ with $s + r \leq \min\{N_1, N_2\}$, and a vector $\mathbf{e} \in \mathbb{R}^m$, if $\text{rank}(\mathbf{X}) \leq r$ and if the modified rank-restricted isometry property (5) holds with ratio $\gamma_{s+r} = \beta_{s+r}/\alpha_{s+r}$, then

$$(10) \quad \|\mathbf{X} - \mu H_s(\mathcal{A}^* \text{sgn}(\mathcal{A}\mathbf{X} + \mathbf{e}))\|_F \leq \left(1 - \frac{1}{2\gamma_{s+r}^2} + \tau \sqrt{\frac{r}{s+r}}\right) \|\mathbf{X}\|_F + \frac{6\tau^2}{\beta_{s+r}} \|\mathbf{e}\|_1,$$

provided the stepsize μ satisfies, for some $\tau \geq 1$,

$$(11) \quad \mu = \frac{\|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1}{\nu} \quad \text{with} \quad \frac{\beta_{s+r}^2}{\tau} \leq \nu \leq \beta_{s+r}^2 \quad \text{and} \quad \nu \geq \|H_s(\mathcal{A}^* \text{sgn}(\mathcal{A}\mathbf{X} + \mathbf{e}))\|_F^2.$$

Propositions 2 and 3 both rely on a purely matrix-analytic observation stated below.

Lemma 4. Given matrices $\mathbf{A}, \mathbf{B} \in \mathbb{R}^{N_1 \times N_2}$ and integers i, j, k with $i \leq k$ and $k + j \leq \min\{N_1, N_2\}$, if $\text{rank}(\mathbf{A}) \leq j$, then

$$(12) \quad |\langle \mathbf{A}, \mathbf{B} - H_k(\mathbf{B}) \rangle_F| \leq \sqrt{\frac{j}{k + j - i}} \|\mathbf{A}\|_F \|H_{k+j}(\mathbf{B}) - H_i(\mathbf{B})\|_F.$$

⁴The expression of $\eta(\kappa)$ in (9) is certainly not optimal, but it suffices for our purpose.

⁵[7, Lemma 3] is not directly exploitable due to the obstruction described in footnote 3.

Proof. We apply Von Neuman's trace inequality, take into account the facts that $\sigma_\ell(\mathbf{A}) = 0$ for $\ell > j$ and that $\sigma_\ell(\mathbf{B} - H_k(\mathbf{B})) = \sigma_{k+\ell}(\mathbf{B})$, and finally use Cauchy–Schwarz inequality to write

$$\begin{aligned}
 (13) \quad |\langle \mathbf{A}, \mathbf{B} - H_k(\mathbf{B}) \rangle_F| &\leq \sum_{\ell \geq 1} \sigma_\ell(\mathbf{A}) \sigma_\ell(\mathbf{B} - H_k(\mathbf{B})) \\
 &= \sum_{\ell=1}^j \sigma_\ell(\mathbf{A}) \sigma_{k+\ell}(\mathbf{B}) \\
 &\leq \left[\sum_{\ell=1}^j \sigma_\ell(\mathbf{A})^2 \right]^{1/2} \left[\sum_{\ell=1}^j \sigma_{k+\ell}(\mathbf{B})^2 \right]^{1/2}.
 \end{aligned}$$

For $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_N$, $N := \min\{N_1, N_2\}$, it is easy to verify that

$$(14) \quad \sum_{\ell=1}^j \lambda_{k+\ell} \leq \frac{j}{k+j-i} \sum_{\ell=1}^{k+j-i} \lambda_{i+\ell}.$$

By applying this inequality with $\lambda_h = \sigma_h(\mathbf{B})^2$ and substituting into (13), we obtain

$$\begin{aligned}
 (15) \quad |\langle \mathbf{A}, \mathbf{B} - H_k(\mathbf{B}) \rangle_F| &\leq \|\mathbf{A}\|_F \left[\frac{j}{k+j-i} \sum_{\ell=1}^{k+j-i} \sigma_{i+\ell}(\mathbf{B})^2 \right]^{1/2} \\
 &= \sqrt{\frac{j}{k+j-i}} \|\mathbf{A}\|_F \|H_{k+j}(\mathbf{B}) - H_i(\mathbf{B})\|_F,
 \end{aligned}$$

which is the announced result. \square

We are now in a position to fully justify Propositions 2 and 3.

Proof of Proposition 2. We start by noticing that

$$\begin{aligned}
 (16) \quad \|\mathbf{X} - H_s(\mathbf{Z})\|_F^2 &= \langle \mathbf{X} - H_s(\mathbf{Z}), \mathbf{X} - \mathbf{Z} \rangle_F + \langle \mathbf{X} - H_s(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F \\
 &\leq \frac{1}{2} \|\mathbf{X} - H_s(\mathbf{Z})\|_F^2 + \frac{1}{2} \|\mathbf{X} - \mathbf{Z}\|_F^2 + \langle \mathbf{X} - H_s(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F.
 \end{aligned}$$

Thus, after rearranging the terms, we have

$$\begin{aligned}
 (17) \quad \|\mathbf{X} - H_s(\mathbf{Z})\|_F^2 &\leq \|\mathbf{X} - \mathbf{Z}\|_F^2 + 2\langle \mathbf{X} - H_s(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F \\
 &= \|\mathbf{X} - \mathbf{Z}\|_F^2 + 2\langle \mathbf{X} - H_r(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F,
 \end{aligned}$$

where the last equality followed from the fact that $\langle H_\ell(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F = 0$ whenever $\ell \leq s$. We apply Lemma 4 with $\mathbf{A} = \mathbf{X} - H_r(\mathbf{Z})$, $\mathbf{B} = \mathbf{Z}$, $i = r$, $j = 2r$, and $k = s$. This allows us to write

$$\begin{aligned}
 (18) \quad \langle \mathbf{X} - H_r(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F &\leq \sqrt{\frac{2r}{s+r}} \|\mathbf{X} - H_r(\mathbf{Z})\|_F \|H_{s+2r}(\mathbf{Z}) - H_r(\mathbf{Z})\|_F \\
 &\leq \sqrt{\frac{2r}{s+r}} \|\mathbf{X} - H_r(\mathbf{Z})\|_F \|\mathbf{Z} - H_r(\mathbf{Z})\|_F.
 \end{aligned}$$

In view of $\|\mathbf{Z} - H_r(\mathbf{Z})\|_F \leq \|\mathbf{Z} - \mathbf{X}\|_F$ (recall that $H_r(\mathbf{Z})$ is a best approximant of rank at most r to \mathbf{Z}) and of $\|\mathbf{X} - H_r(\mathbf{Z})\|_F \leq \|\mathbf{Z} - H_r(\mathbf{Z})\|_F + \|\mathbf{Z} - \mathbf{X}\|_F \leq 2\|\mathbf{Z} - \mathbf{X}\|_F$, we now arrive at

$$(19) \quad \langle \mathbf{X} - H_r(\mathbf{Z}), \mathbf{Z} - H_s(\mathbf{Z}) \rangle_F \leq 2\sqrt{\frac{2r}{s+r}} \|\mathbf{Z} - \mathbf{X}\|_F^2.$$

Substituting (19) into (17) leads to

$$(20) \quad \|\mathbf{X} - H_s(\mathbf{Z})\|_F^2 \leq \left(1 + 4\sqrt{\frac{2r}{s+r}}\right) \|\mathbf{Z} - \mathbf{X}\|_F^2 \leq \left(1 + 2\sqrt{\frac{2r}{s+r}}\right)^2 \|\mathbf{X} - \mathbf{Z}\|_F^2,$$

which yields the desired result after taking the square root on both sides. \square

Proof of Proposition 3. For convenience, we denote $\mathbf{w} := \text{sgn}(\mathcal{A}\mathbf{X} + \mathbf{e})$ throughout the proof. Before addressing the bound (10) on $\|\mathbf{X} - \mu H_s(\mathcal{A}^*\mathbf{w})\|_F$, we remark that $\|H_s(\mathcal{A}^*\mathbf{w})\|_F^2 \leq \beta_{s+r}^2$, so that (11) does not feature a vacuous requirement. In fact, a stronger inequality holds, namely

$$(21) \quad \|H_{s+r}(\mathcal{A}^*\mathbf{w})\|_F \leq \beta_{s+r}.$$

This follows from the observation that

$$(22) \quad \begin{aligned} \|H_{s+r}(\mathcal{A}^*\mathbf{w})\|_F^2 &= \langle H_{s+r}(\mathcal{A}^*\mathbf{w}), H_{s+r}(\mathcal{A}^*\mathbf{w}) \rangle_F = \langle \mathcal{A}^*\mathbf{w}, H_{s+r}(\mathcal{A}^*\mathbf{w}) \rangle_F \\ &= \langle \mathbf{w}, \mathcal{A}(H_{s+r}(\mathcal{A}^*\mathbf{w})) \rangle \leq \|\mathcal{A}(H_{s+r}(\mathcal{A}^*\mathbf{w}))\|_1 \\ &\leq \beta_{s+r} \|H_{s+r}(\mathcal{A}^*\mathbf{w})\|_F. \end{aligned}$$

Turning now to $\|\mathbf{X} - \mu H_s(\mathcal{A}^*\mathbf{w})\|_F$, we expand its square to obtain

$$(23) \quad \begin{aligned} \|\mathbf{X} - \mu H_s(\mathcal{A}^*\mathbf{w})\|_F^2 &= \|\mathbf{X}\|_F^2 - 2\mu \langle \mathbf{X}, H_s(\mathcal{A}^*\mathbf{w}) \rangle_F + \mu^2 \|H_s(\mathcal{A}^*\mathbf{w})\|_F^2 \\ &\leq \|\mathbf{X}\|_F^2 - 2\mu \langle \mathbf{X}, H_s(\mathcal{A}^*\mathbf{w}) \rangle_F + \mu^2 \nu. \end{aligned}$$

To deal with the inner product appearing on the right-hand side of (23), we write

$$(24) \quad \langle \mathbf{X}, H_s(\mathcal{A}^*\mathbf{w}) \rangle_F = \langle \mathbf{X}, \mathcal{A}^*\mathbf{w} \rangle_F - \langle \mathbf{X}, \mathcal{A}^*\mathbf{w} - H_s(\mathcal{A}^*\mathbf{w}) \rangle_F,$$

while noticing that

$$(25) \quad \langle \mathbf{X}, \mathcal{A}^*\mathbf{w} \rangle_F = \langle \mathcal{A}\mathbf{X}, \text{sgn}(\mathcal{A}\mathbf{X} + \mathbf{e}) \rangle = \|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1 - \langle \mathbf{e}, \text{sgn}(\mathcal{A}\mathbf{X} + \mathbf{e}) \rangle \geq \|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1 - \|\mathbf{e}\|_1$$

and that Lemma 4 applied with $\mathbf{A} = \mathbf{X}$, $\mathbf{B} = \mathcal{A}^*\mathbf{w}$, $i = 0$, $j = r$, and $k = s$ yields

$$(26) \quad |\langle \mathbf{X}, \mathcal{A}^*\mathbf{w} - H_s(\mathcal{A}^*\mathbf{w}) \rangle_F| \leq \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F \|H_{s+r}(\mathcal{A}^*\mathbf{w})\|_F.$$

Taking (21) into account and making use of (25) and (26) in (24), we derive

$$(27) \quad \langle \mathbf{X}, H_s(\mathcal{A}^*\mathbf{w}) \rangle_F \geq \|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1 - \left[\|\mathbf{e}\|_1 + \beta_{s+r} \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F \right].$$

Substituting (27) into (23) and using (11) now gives

$$\begin{aligned}
 (28) \quad \|\mathbf{X} - \mu H_s(\mathcal{A}^* \mathbf{w})\|_F^2 &\leq \|\mathbf{X}\|_F^2 - 2\mu \|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1 + \mu^2 \nu + 2\mu \left[\|\mathbf{e}\|_1 + \beta_{s+r} \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F \right] \\
 &= \|\mathbf{X}\|_F^2 - \frac{\|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1^2}{\nu} + 2\mu \left[\|\mathbf{e}\|_1 + \beta_{s+r} \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F \right] \\
 &\leq \|\mathbf{X}\|_F^2 - \frac{\|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1^2}{\beta_{s+r}^2} + \frac{2\tau \|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1}{\beta_{s+r}^2} \left[\|\mathbf{e}\|_1 + \beta_{s+r} \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F \right].
 \end{aligned}$$

In view of the inequalities $\|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1^2 \geq \| \|\mathcal{A}\mathbf{X}\|_1 - \|\mathbf{e}\|_1 \|^2 = \|\mathcal{A}\mathbf{X}\|_1^2 - 2\|\mathcal{A}\mathbf{X}\|_1 \|\mathbf{e}\|_1 + \|\mathbf{e}\|_1^2$ and $\|\mathcal{A}\mathbf{X} + \mathbf{e}\|_1 \leq \|\mathcal{A}\mathbf{X}\|_1 + \|\mathbf{e}\|_1$, we arrive at

$$(29) \quad \|\mathbf{X} - \mu H_s(\mathcal{A}^* \mathbf{w})\|_F^2 \leq a + b \|\mathbf{e}\|_1 + c \|\mathbf{e}\|_1^2, \quad \begin{cases} a := \|\mathbf{X}\|_F^2 - \frac{\|\mathcal{A}\mathbf{X}\|_1^2}{\beta_{s+r}^2} + \frac{2\tau}{\beta_{s+r}} \sqrt{\frac{r}{s+r}} \|\mathcal{A}\mathbf{X}\|_1 \|\mathbf{X}\|_F, \\ b := \frac{2(1+\tau) \|\mathcal{A}\mathbf{X}\|_1}{\beta_{s+r}^2} + \frac{2\tau}{\beta_{s+r}} \sqrt{\frac{r}{s+r}} \|\mathbf{X}\|_F, \\ c := \frac{2\tau - 1}{\beta_{s+r}^2}. \end{cases}$$

The bounds $\|\mathcal{A}\mathbf{X}\|_1 \geq \alpha_{s+r} \|\mathbf{X}\|_F$ and $\|\mathcal{A}\mathbf{X}\|_1 \leq \beta_{s+r} \|\mathbf{X}\|_F$ imply that

$$(30) \quad \begin{cases} a \leq \left(1 - \frac{1}{\gamma_{s+r}^2} + 2\tau \sqrt{\frac{r}{s+r}}\right) \|\mathbf{X}\|_F^2 \leq \left(1 - \frac{1}{2\gamma_{s+r}^2} + \tau \sqrt{\frac{r}{s+r}}\right)^2 \|\mathbf{X}\|_F^2, \\ b \leq \frac{2}{\beta_{s+r}} \left(1 + \tau + \tau \sqrt{\frac{r}{s+r}}\right) \|\mathbf{X}\|_F \leq \frac{2\tau}{\beta_{s+r}} \left(1 + \tau + \tau \sqrt{\frac{r}{s+r}}\right) \|\mathbf{X}\|_F. \end{cases}$$

Since we also have

$$(31) \quad c \leq \frac{\tau^2}{\beta_{s+r}^2} \leq \frac{\tau^2 \left(1 + \tau + \tau \sqrt{\frac{r}{s+r}}\right)^2}{\beta_{s+r}^2 \left(1 - \frac{1}{2\gamma_{s+r}^2} + \tau \sqrt{\frac{r}{s+r}}\right)^2},$$

we deduce that $\|\mathbf{X} - \mu H_s(\mathcal{A}^* \mathbf{w})\|_F^2$ is bounded above by

$$(32) \quad \left[\left(1 - \frac{1}{2\gamma_{s+r}^2} + \tau \sqrt{\frac{r}{s+r}}\right) \|\mathbf{X}\|_F + \frac{\tau \left(1 + \tau + \tau \sqrt{\frac{r}{s+r}}\right)}{\beta_{s+r} \left(1 - \frac{1}{2\gamma_{s+r}^2} + \tau \sqrt{\frac{r}{s+r}}\right)} \|\mathbf{e}\|_1 \right]^2.$$

It remains to take the square root on both sides to obtain the desired result, after noticing that $1 + \tau + \tau \sqrt{r/(s+r)} \leq 3\tau$ and $1 - 1/(2\gamma_{s+r}^2) + \tau \sqrt{r/(s+r)} \geq 1 - 1/(2\gamma_{s+r}^2) \geq 1/2$. \square

3 Stability and robustness of the reconstruction

This section is devoted to the proof of Theorem 1 — in fact, of the more general theorem stated below, which incorporates low-rank defect and measurement error.

Theorem 5. Let $\mathcal{A} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^m$ be a measurement map satisfying the modified rank-restricted isometry property (5) of order $\gamma_{t+s+r} \leq \gamma$, where

$$(33) \quad s = 100\gamma^4 r \quad \text{and} \quad t = 800\gamma^{12} s.$$

For all $\mathbf{X} \in \mathbb{R}^{N_1 \times N_2}$ and all $\mathbf{e} \in \mathbb{R}^m$, the iterative hard thresholding algorithm (8) applied to $\mathbf{y} = \mathcal{A}\mathbf{X} + \mathbf{e}$ with stepsize

$$(34) \quad \mu_n = \frac{\|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1}{\max \left\{ \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2, \frac{1}{4\gamma^2} \frac{\|\mathcal{A}(H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))\|_1^2}{\|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2} \right\}},$$

produces a sequence $(\mathbf{X}_n)_{n \geq 0}$ whose cluster points \mathbf{X}^\sharp approximate $H_r(\mathbf{X})$ with error

$$(35) \quad \|H_r(\mathbf{X}) - \mathbf{X}^\sharp\|_F \leq d \|\mathcal{A}(\mathbf{X} - H_r(\mathbf{X})) + \mathbf{e}\|_1,$$

where the constant d depends only on γ and β_{t+s+r} .

A number of comments are on order before giving the full justification of Theorem 5.

The parameters s and t . The constants involved in (33) are admittedly huge. They have been chosen to look nice and to make the theory work, but they do not reflect reality. In practice, the numerical experiments carried out in Section 4 suggest that they can be taken much smaller.

The idealized situation. Theorem 1 is truly a special case of Theorem 5 where one assumes that $\text{rank}(\mathbf{X}) \leq r$ and $\mathbf{e} = \mathbf{0}$. Not only does (35) guarantee exact recovery in this case, but the stepsize (34) indeed reduces to $\mu_n = \|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1 / \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2$ when $\text{rank}(\mathbf{X}) \leq r$ and $\mathbf{e} = \mathbf{0}$. This is valid because, on the one hand,

$$(36) \quad \frac{1}{4\gamma^2} \frac{\|\mathcal{A}(H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))\|_1^2}{\|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2} \leq \frac{1}{4\gamma^2} \beta_t^2 \leq \frac{1}{4\gamma^2} \beta_{t+s+r}^2 \leq \frac{\alpha_{t+s+r}^2}{4},$$

and, on the other hand, $\|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2 \geq \|H_{s+r}(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2 \geq \alpha_{t+s+r}^2$, since

$$(37) \quad \begin{aligned} \alpha_{t+s+r} \|\mathbf{X} - \mathbf{X}_n\|_F &\leq \|\mathcal{A}(\mathbf{X} - \mathbf{X}_n)\|_1 = \langle \mathcal{A}(\mathbf{X} - \mathbf{X}_n), \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n)) \rangle \\ &= \langle \mathbf{X} - \mathbf{X}_n, \mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n) \rangle \leq \sum_{\ell=1}^{s+r} \sigma_\ell(\mathbf{X} - \mathbf{X}_n) \sigma_\ell(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)) \\ &\leq \|\mathbf{X} - \mathbf{X}_n\|_F \|H_{s+r}(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F. \end{aligned}$$

The stepsize. The selection made in (34) is certainly not the only possible one. For instance, choosing $\mu_n = \|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1 / \beta_{t+s+r}^2$ would also work. We favored the option (34) because it does not necessitate an estimation of β_{t+s+r} , because it seems to perform better in practice, and because it is close to optimal in the sense that the choice $\mu_n = \|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1 / \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2$, made in the idealized situation, almost minimizes the expression (23) when taking (27) into account. Let us also remark that there is some leeway around an admissible μ_n , namely replacing it by $\mu'_n = (1 + \delta_n)\mu_n$ for a suitably small relative variation δ_n does not compromise the end result, as it only creates a minor change in the proof (specifically, in (57)). The reader is invited to fill in the details.

Number of iterations. The proof of Theorem 5 actually provides error estimates throughout the iterations. For instance, in the idealized situation where $\text{rank}(\mathbf{X}) \leq r$ and $\mathbf{e} = \mathbf{0}$, it reveals that there is a constant $\rho < 1$ such that

$$(38) \quad \|\mathbf{X} - \mathbf{X}_n\|_F \leq \rho^n \|\mathbf{X}\|_F \quad \text{for all } n \geq 0.$$

Thus, for a prescribed accuracy ε , we have $\|\mathbf{X} - \mathbf{X}_n\|_F \leq \varepsilon$ in only $n = \lceil \ln(\|\mathbf{X}\|_F / \varepsilon) / \ln(1/\rho) \rceil$ iterations.

A stopping criterion. Although Theorem 5 states a recovery guarantee for cluster points of the sequence $(\mathbf{X}_n)_{n \geq 0}$, one can stop the iterative process as soon as the maximum in (34) is achieved by the second term defining it (provided this occurs). Precisely, if

$$(39) \quad \|\mathcal{A}(H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))\|_1 \geq 2\gamma \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2,$$

then the conclusion that

$$(40) \quad \|H_r(\mathbf{X}) - \mathbf{X}_n\|_F \leq d \|\mathbf{e}'\|_1, \quad \mathbf{e}' := \mathcal{A}(\mathbf{X} - H_r(\mathbf{X})) + \mathbf{e}$$

is already acquired. Indeed, the condition (39) imposes that

$$(41) \quad 2\gamma \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F^2 \leq \beta_{t+s+r} \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F,$$

i.e., that

$$(42) \quad \|H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F \leq \frac{\beta_{t+s+r}}{2\gamma}.$$

We then derive (as in (37)) that

$$(43) \quad \begin{aligned} \langle H_r(\mathbf{X}) - \mathbf{X}_n, \mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n) \rangle &\leq \|H_r(\mathbf{X}) - \mathbf{X}_n\|_F \|H_{s+r}(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))\|_F \\ &\leq \frac{\beta_{t+s+r}}{2\gamma} \|H_r(\mathbf{X}) - \mathbf{X}_n\|_F. \end{aligned}$$

Moreover, keeping on mind that $\mathbf{y} = \mathcal{A}(H_r(\mathbf{X})) + \mathbf{e}'$, we notice that

$$\begin{aligned}
 (44) \quad \langle H_r(\mathbf{X}) - \mathbf{X}_n, \mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n) \rangle &= \langle \mathcal{A}(H_r(\mathbf{X}) - \mathbf{X}_n), \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n) \rangle \\
 &= \langle \mathbf{y} - \mathcal{A}\mathbf{X}_n - \mathbf{e}', \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n) \rangle \geq \|\mathbf{y} - \mathcal{A}\mathbf{X}_n\|_1 - \|\mathbf{e}'\|_1 \\
 &\geq \|\mathcal{A}(H_r(\mathbf{X}) - \mathbf{X}_n)\|_1 - 2\|\mathbf{e}'\|_1 \geq \alpha_{t+s+r} \|H_r(\mathbf{X}) - \mathbf{X}_n\|_F - 2\|\mathbf{e}'\|_1 \\
 &\geq \frac{\beta_{t+s+r}}{\gamma} \|H_r(\mathbf{X}) - \mathbf{X}_n\|_F - 2\|\mathbf{e}'\|_1.
 \end{aligned}$$

Combining (43) and (44) yields

$$(45) \quad \|H_r(\mathbf{X}) - \mathbf{X}_n\|_1 \leq \frac{4\gamma}{\beta_{t+s+r}} \|\mathbf{e}'\|_1.$$

Form of the error estimate. In the realistic situation where \mathbf{X} is not exactly low-rank and measurement error occurs, one customarily sees recovery guarantees of the type

$$(46) \quad \|\mathbf{X} - \mathbf{X}^\sharp\|_F \leq \frac{C}{\sqrt{r}} \min_{\text{rank}(\mathbf{Z}) \leq r} \|\mathbf{X} - \mathbf{Z}\|_* + D\|\mathbf{e}\|_1.$$

Our iterative hard thresholding algorithm (with parameter r replaced by $2r$) also allows for such an estimate, which is derived using the sort-and-split technique. We omit the details, as the argument is quite classical and can be found in full in [6] for the vector case.

The norm on the measurement error. The appearance of $\|\mathbf{e}\|_1$ in (46) might seem surprising, as one usually encounters $\|\mathbf{e}\|_2$ when the measurement maps satisfies the standard rank-restricted isometry property. But there is no discrepancy here, as this is just a matter of normalization. Indeed, for a measurement map $\mathcal{G} : \mathbb{R}^{N_1 \times N_2} \rightarrow \mathbb{R}^m$ generated as in (7) with independent standard normal random variables $\mathcal{G}_{i,(k,\ell)}$, let us suppose that inaccurate measurements are given by $\mathbf{y} = \mathcal{G}\mathbf{X} + \mathbf{u}$. On the one hand, the normalized version of \mathcal{G} that satisfies the standard rank-restricted isometry property is $\mathcal{A} = \mathcal{G}/\sqrt{m}$, in which case the measurements take the form $\mathbf{y}/\sqrt{m} = \mathcal{A}\mathbf{X} + \mathbf{e}$ with $\mathbf{e} = \mathbf{u}/\sqrt{m}$, so that $\|\mathbf{e}\|_2 = \|\mathbf{u}\|_2/\sqrt{m}$ appears in the error estimate. On the other hand, the normalized version of \mathcal{G} that satisfies the modified rank-restricted isometry property is $\mathcal{A} = \mathcal{G}/m$, in which case the measurements take the form $\mathbf{y}/m = \mathcal{A}\mathbf{X} + \mathbf{e}$ with $\mathbf{e} = \mathbf{u}/m$, so that $\|\mathbf{e}\|_1 = \|\mathbf{u}\|_1/m$ appears in the error estimate. The latter is in fact better than the former, by virtue of $\|\mathbf{u}\|_1 \leq \sqrt{m}\|\mathbf{u}\|_2$ for $\mathbf{u} \in \mathbb{R}^m$.

Proof of Theorem 5. Without loss of generality, we can assume that \mathbf{X} has rank at most r and aim at proving that

$$(47) \quad \|\mathbf{X} - \mathbf{X}^\sharp\|_F \leq d\|\mathbf{e}\|_1.$$

Indeed, we can in general interpret the measurements $\mathbf{y} = \mathcal{A}\mathbf{X} + \mathbf{e}$ made on \mathbf{X} as measurements $\mathbf{y} = \mathcal{A}(H_r(\mathbf{X})) + \mathbf{e}'$, $\mathbf{e}' := \mathcal{A}(\mathbf{X} - H_r(\mathbf{X})) + \mathbf{e}$, made on the matrix $H_r(\mathbf{X})$ which has rank is at

most r . The exact low-rank result would then yield

$$(48) \quad \|H_r(\mathbf{X}) - \mathbf{X}^\sharp\|_F \leq d\|\mathbf{e}'\|_1 = d\|\mathcal{A}(\mathbf{X} - H_r(\mathbf{X})) + \mathbf{e}\|_1,$$

as desired. So we assume from now on that $\text{rank}(\mathbf{X}) \leq r$. We are going to establish the existence of $\rho < 1$ such that, for all $n \geq 0$,

$$(49) \quad \|\mathbf{X} - \mathbf{X}_{n+1}\|_F \leq \rho\|\mathbf{X} - \mathbf{X}_n\|_F + d'\|\mathbf{e}\|_1,$$

which immediately implies that, for all $n \geq 0$,

$$(50) \quad \|\mathbf{X} - \mathbf{X}_n\|_F \leq \rho^n\|\mathbf{X}\|_F + d\|\mathbf{e}\|_1, \quad d := \frac{d'}{1-\rho}.$$

In particular, (47) is obtained from (50) by letting $n \rightarrow \infty$. It remains to prove (49) for a given integer $n \geq 0$. Since $\mathbf{X}_{n+1} = H_s(\mathbf{X}_n + \mu_n H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))$, Proposition 2 implies

$$(51) \quad \|\mathbf{X} - \mathbf{X}_{n+1}\|_F \leq \rho'\|\mathbf{X} - (\mathbf{X}_n + \mu_n H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))\|_F, \quad \rho' := 1 + \sqrt{\frac{8r}{s+r}} > 1.$$

Next, we notice that the stepsize μ_n takes the form

$$(52) \quad \mu_n = \frac{\|\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}\|_1}{\nu_n},$$

where the denominator

$$(53) \quad \nu_n := \max \left\{ \|H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F^2, \frac{1}{4\gamma^2} \frac{\|\mathcal{A}(H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e})))\|_1^2}{\|H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F^2} \right\}$$

clearly satisfies

$$(54) \quad \nu_n \geq \|H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F^2,$$

as well as

$$(55) \quad \frac{\beta_{t+s+r}^2}{4\gamma^4} \leq \nu_n \leq \beta_{t+s+r}^2.$$

The latter inequalities follow from $\|H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F^2 \leq \beta_t^2 \leq \beta_{t+s+r}^2$ (obtained in the same way as (21)) and from

$$(56) \quad \frac{1}{4\gamma^2} \frac{\|\mathcal{A}(H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e})))\|_1^2}{\|H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F^2} \begin{cases} \leq \frac{1}{4\gamma^2} \beta_{t+s+r}^2 \leq \beta_{t+s+r}^2, \\ \geq \frac{1}{4\gamma^2} \alpha_{t+s+r}^2 \geq \frac{\beta_{t+s+r}^2}{4\gamma^4}. \end{cases}$$

Therefore, according to Proposition 3, we obtain

$$(57) \quad \begin{aligned} \|\mathbf{X} - (\mathbf{X}_n + \mu_n H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n)))\|_F &= \|(\mathbf{X} - \mathbf{X}_n) - \mu_n H_t(\mathcal{A}^* \text{sgn}(\mathcal{A}(\mathbf{X} - \mathbf{X}_n) + \mathbf{e}))\|_F \\ &\leq \rho''\|\mathbf{X} - \mathbf{X}_n\|_F + \frac{96\gamma^8}{\beta_{t+s+r}}\|\mathbf{e}\|_1, \quad \rho'' := 1 - \frac{1}{2\gamma^2} + 4\gamma^4 \sqrt{\frac{s+r}{t+s+r}} < 1, \end{aligned}$$

with $4\gamma^4\sqrt{(s+r)/(t+s+r)} < 1/(2\gamma^2)$ being justified by our choice of s and t . In fact, our choice of s and t guarantees that

$$(58) \quad \begin{aligned} \rho := \rho'\rho'' &\leq \left(1 + \sqrt{\frac{8r}{s}}\right) \left(1 - \frac{1}{2\gamma^2} + 4\gamma^4\sqrt{\frac{2s}{t}}\right) \leq 1 - \frac{1}{2\gamma^2} + 4\gamma^4\sqrt{\frac{2s}{t}} + \sqrt{\frac{8r}{s}} \\ &\leq 1 - \frac{1}{2\gamma^2} + \frac{1}{5\gamma^2} + \frac{\sqrt{2}}{5\gamma^2} = 1 - \frac{3 - 2\sqrt{2}}{10\gamma^2} < 1. \end{aligned}$$

Combining (51) and (57), we finally derive that

$$(59) \quad \|\mathbf{X} - \mathbf{X}_{n+1}\|_F \leq \rho\|\mathbf{X} - \mathbf{X}_n\|_F + d'\|\mathbf{e}\|_1, \quad d' := \frac{96\rho'\gamma^8}{\beta_{t+s+r}}.$$

This establishes (49) and completes the proof. \square

4 Numerical validation

The perspective adopted in this article is mostly theoretical, with a declared goal consisting in uncovering iterative thresholding algorithms that provide alternatives to nuclear norm minimization when the standard rank-restricted isometry property does not hold. The modest experiments carried out with our proposed algorithm are encouraging, although we shall refrain from making too optimistic claims about its performance. In what follows, we only focus on measurements obtained by Gaussian rank-one projections of type (6). Note that in this case our algorithm starts with $\mathbf{X}_0 = \mathbf{0}$ and iterates the scheme

$$(60) \quad \mathbf{X}_{n+1} = H_s \left(\mathbf{X}_n + \mu_n H_t \left(\sum_{i=1}^m \text{sgn}(y_i - \mathbf{a}_i^* \mathbf{X}_n \mathbf{b}_i) \mathbf{b}_i \mathbf{a}_i^* \right) \right).$$

All the experiments performed below can be reproduced by downloading the MATLAB file linked to this article from the first author's webpage.

Influence of s and t . In order to derive theoretical guarantees about the success of the algorithm, the parameters s and t appearing in (60) can be chosen according to (33). This of course does not mean that other choices are unsuitable. Figure 1 reports on the effect of varying these parameters. We observe that successful recoveries occur more frequently when $s = r$, which makes intuitive sense since the rank- s matrix \mathbf{X}_n is supposed to approximate the rank- r matrix \mathbf{X} . We also observe that successful recoveries occur more frequently when t increases. The largest value t can take is $N = \min\{N_1, N_2\}$, in which case the hard thresholding operator does not do anything. The default value of the parameters are therefore $s = r$ and $t = N$, meaning that H_t is absent and that one singular value decomposition is saved per iteration.




Figure 1: Influence of s and t on the success of algorithm (60).

Performance. We now compare the low-rank recovery capability of the algorithm (60) with two other algorithms, namely with nuclear norm minimization (NNM) and with the normalized iterative hard thresholding (NIHT) algorithm proposed in [13]. In fact, we test two versions of our modified iterative hard thresholding (MIHT) algorithm: one with the default parameters as discussed above, and one with parameters $s = r$ and $t = 2r$ — this is another natural choice, as $H_t(\mathcal{A}^* \text{sgn}(\mathbf{y} - \mathcal{A}\mathbf{X}_n))$ is supposed to approximate the rank- $2r$ matrix $\mathbf{X} - \mathbf{X}_n$. Judging from Figure 2, NNM is seemingly (and surprisingly) outperformed by our default MIHT algorithm. The comparison with NIHT is less straightforward. With few measurements available, NIHT does best, but when the number of measurements is large enough for the recoveries by NNM and MIHT to be certain, recovery by NIHT does not occur consistently. We interpret this phenomenon as a reflection of the theory. Indeed, Gaussian rank-one projections of type (6) do not obey the standard rank-restricted isometry property, which would have guaranteed the success of NIHT (see [13]), but they obey the modified rank-restricted isometry property (5), which does guarantee the success of NNM (see [1]) and of MIHT (this article).

Speed. The next experiment compares the four algorithms previously considered in terms of execution time and number of iterations (when applicable). Figure 3 shows, as expected, that iterative algorithms are much faster than nuclear norm minimization: in this experiment (which only takes successful recoveries into account), the execution time is $\Theta(N^\omega)$ with $\omega \approx 4$ for NNM, $\omega \approx 1$ for NIHT, and $\omega \approx 2.4$ for MIHT. The execution times for the two versions of MIHT essentially differ by a factor of two, due to the fact that the default version performs only one singular value decomposition per iteration instead of two. The number of iterations is roughly similar for the two versions. It is much lower for NIHT, which explains its superiority in terms of overall execution time.




Figure 2: Frequency of successful recoveries as a function of the number of Gaussian rank-one measurements for nuclear norm minimization, normalized iterative hard thresholding, and modified iterative hard thresholding with default parameters and with parameters $s = r$ and $t = 2r$.

Robustness. The final experiment is designed to examine the effect on the recovery of errors $\mathbf{e} \in \mathbb{R}^m$ in the measurements $\mathbf{y} = \mathcal{A}(\mathbf{X}) + \mathbf{e}$ made on low-rank matrices \mathbf{X} . In particular, we wish to certify the statement (35) that the recovery error, measured in Frobenius norm, is at most proportional to the measurement error, measured in ℓ_1 -norm. The results displayed in Figure 4 confirm the validity of this statement for our algorithm (60) with default parameters and with parameters $s = r$ and $t = 2r$, as well as for the normalized iterative hard thresholding algorithm.

5 Explanation of the modified rank-restricted isometry property

In this appendix, we point out that the modified rank-restricted isometry property (5) is not only valid for rank-one Gaussian measurements, as already established in [1], but also for subexponential measurements of type (7). The justification essentially follows the arguments put forward in [5]. First, fixing a matrix $\mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}$, viewed as its vectorization $\mathbf{z} = \text{vec}(\mathbf{Z}) \in \mathbb{R}^{N_1 N_2}$, [5, Theorem 3.1] guarantees a concentration inequality of the form

$$(61) \quad \Pr(\left| \|\mathcal{A}(\mathbf{Z})\|_1 - \|\mathbf{Z}\| \right| > \varepsilon \|\mathbf{Z}\|) \leq 2 \exp(-\kappa \varepsilon^2 m).$$

The constant κ depends on the subexponential distribution, and so does the slanted norm, although it is comparable to the Frobenius norm in the sense that (see [5, Proposition 2.3])

$$(62) \quad c \|\mathbf{Z}\|_F \leq \|\mathbf{Z}\| \leq C \|\mathbf{Z}\|_F \quad \text{for all } \mathbf{Z} \in \mathbb{R}^{N_1 \times N_2}$$

for some constants $c, C > 0$ depending on the subexponential distribution. The next step (see the proof of [5, Theorem 4.1]) is a covering argument showing that (61) extends to a uniform




Figure 3: Execution time and number of iterations as a function of the dimension $N = N_1 = N_2$.

concentration inequality of the form

$$(63) \quad \Pr(\|\mathcal{A}(\mathbf{Z})\|_1 - \|\mathbf{Z}\| > \delta \|\mathbf{Z}\| \text{ for some } \mathbf{Z} \text{ with } \text{rank}(\mathbf{Z}) \leq r) \leq 2 \exp(-\kappa \bar{\delta}^2 m + \ln(\mathcal{N}_r(\bar{\delta}))),$$

where $\mathcal{N}_r(\bar{\delta})$, $\bar{\delta} := \delta/3$, denotes the size of a $\bar{\delta}$ -net for the set of rank- r matrices with slanted norm at most 1. We now recall the key observation from [2] that the set of rank- r matrices in $\mathbb{R}^{N_1 \times N_2}$ with norm at most 1 has a covering number $\mathcal{N}_r(\bar{\delta}) \leq (c_1/\bar{\delta})^{c_2 r} N_{\max}$, $N_{\max} := \max\{N_1, N_2\}$, which is valid for the Frobenius norm and, in view of (62), for the slanted norm, too. This implies a modified rank-restricted isometry property of the form

$$(64) \quad (1 - \delta) \|\mathbf{Z}\| \leq \|\mathcal{A}(\mathbf{Z})\|_1 \leq (1 + \delta) \|\mathbf{Z}\| \quad \text{whenever } \text{rank}(\mathbf{Z}) \leq r,$$

which occurs with failure probability at most

$$(65) \quad 2 \exp(-\kappa \bar{\delta}^2 m + \ln(c_1/\bar{\delta})^{c_2 r} N_{\max}) \leq 2 \exp(-\kappa \bar{\delta}^2 m/2) \leq 2 \exp(-\kappa' \delta^2 m),$$

provided $\ln(c_1/\bar{\delta})^{c_2 r} N_{\max} \leq \kappa' \bar{\delta}^2 m/2$, i.e., $m \geq C(\delta) r N_{\max}$. Finally, (64) turns into (5) by virtue of (62).

References

- [1] T. Cai and A. Zhang. ROP: Matrix recovery via rank-one projections. *The Annals of Statistics* 43.1 (2015): 102–138.
- [2] E. Candès and Y. Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory* 57.4 (2011): 2342–2359.




Figure 4: Frobenius norm of the recovery error as a function of the ℓ_1 -norm of the measurement error.

- [3] E. Candès, T. Strohmer, and V. Voroninski. PhaseLift: exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics* 66.8 (2013): 1241–1274.
- [4] Y. Chen and E. Candès. Solving random quadratic systems of equations is nearly as easy as solving linear systems. *Communications on Pure and Applied Mathematics* 70.5 (2017): 822–883.
- [5] S. Foucart and M.-J. Lai. Sparse recovery with pre-Gaussian random matrices. *Studia Mathematica* 200.1 (2010): 91–102.
- [6] S. Foucart and G. Lécué. An IHT algorithm for sparse recovery from subexponential measurements. *IEEE Signal Processing Letters* 24.9 (2017): 1280–1283.
- [7] S. Foucart and R. Lynch. Recovering low-rank matrices from binary measurements. Preprint.
- [8] D. Goldfarb and S. Ma. Convergence of fixed-point continuation algorithms for matrix rank minimization. *Foundations of Computational Mathematics* 11.2 (2011): 183–210.
- [9] C. Hegde, P. Indyk, and L. Schmidt. Approximation algorithms for model-based compressive sensing. *IEEE Transactions on Information Theory* 61.9 (2015): 5129–5147.
- [10] P. Jain, R. Meka, and I. Dhillon. Guaranteed rank minimization via singular value projection. *Advances in Neural Information Processing Systems* 23 (NIPS 2010), pp. 937–945.
- [11] B. Recht, M. Fazel, and P. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review* 52.3 (2010): 471–501.

- [12] J. Shen and P. Li. A tight bound of hard thresholding. arXiv:1605.01656 (2016).
- [13] J. Tanner and K. Wei. Normalized iterative hard thresholding for matrix completion. *SIAM Journal on Scientific Computing* 35.5 (2013): S104–S125.