Improved Sparse Low-Rank Matrix Estimation

Ankit Parekh^{a,*}, Ivan W. Selesnick^b

^aDept. of Mathematics, Tandon School of Engineering, New York University

^bDept. of Electrical and Computer Engineering, Tandon School of Engineering, New York University

Abstract

We address the problem of estimating a sparse low-rank matrix from its noisy observation. We propose an objective function consisting of a data-fidelity term and two parameterized non-convex penalty functions. Further, we show how to set the parameters of the non-convex penalty functions, in order to ensure that the objective function is strictly convex. The proposed objective function better estimates sparse low-rank matrices than a convex method which utilizes the sum of the nuclear norm and the ℓ_1 norm. We derive an algorithm (as an instance of ADMM) to solve the proposed problem, and guarantee its convergence provided the scalar augmented Lagrangian parameter is set appropriately. We demonstrate the proposed method for denoising an audio signal and an adjacency matrix representing protein interactions in the ‘Escherichia coli’ bacteria.

Keywords: Low-rank matrix, sparse matrix, speech denoising, nonconvex regularization, convex optimization

1. Introduction

We aim to estimate a sparse low-rank matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ from its noisy observation $\mathbf{Y} \in \mathbb{R}^{m \times n}$, i.e.,

$$\mathbf{Y} = \mathbf{X} + \mathbf{W}, \quad \mathbf{W} \in \mathbb{R}^{m \times n}, \quad (1)$$

where $\mathbf{W}$ represents additive white Gaussian noise (AWGN) matrix. The estimation of sparse low-rank matrices has been studied [7] and used for various applications such as covariance matrix estimation [4, 21, 54], subspace clustering [27], biclustering [34], sparse reduced rank regression [8, 15], graph denoising and link prediction [47, 46], image classification [53] and hyperspectral unmixing [24].

In order to estimate the sparse low-rank matrix $\mathbf{X}$, it has been proposed [47] to solve the following optimization problem

$$\arg \min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \left\{ \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \|\mathbf{X}\|_* + \lambda_1 \|\mathbf{X}\|_1 \right\}, \quad (2)$$

where $\|\cdot\|_*$ is the nuclear norm, $\|\cdot\|_1$ is the entry-wise ℓ_1 norm and $\lambda_i \geq 0$ are the regularization parameters. The nuclear norm induces sparsity of the singular values of the matrix $\mathbf{X}$, while the entry-wise ℓ_1 norm induces sparsity of the elements of $\mathbf{X}$.

The nuclear norm and the ℓ_1 norm are convex relaxations of the non-convex rank and sparsity constraints, respectively. The nuclear norm can be considered as the ℓ_1 norm applied to the singular values of the matrix. It

is known that the ℓ_1 norm underestimates non-zero signal values, when used as a sparsity-inducing regularizer. As a result, the sparse low-rank (SLR) problem in (2) can be considered, in general, to be over-relaxed [32]. Further, it is known that the performance of nuclear norm for sparse regularization of the singular values is sub-optimal [39].

In order to estimate the non-zero signal values more accurately, non-convex regularization has been favored over convex regularization [13, 51, 45, 43, 52]. Furthermore, it has been shown that non-convex penalty functions can induce sparsity of the singular values more effectively than the nuclear norm [37, 26, 28, 14, 40]. Indeed, it was shown that nonconvex regularizers are better able to estimate simultaneously sparse and low-rank matrices in the context of spectral unmixing for hyperspectral images [24]. The use of non-convex regularizers (penalty functions), however, generally leads to non-convex optimization problems. The non-convex optimization problems suffer from numerous issues (sub-optimal local minima, sensitivity to changes in the input data and the regularization parameters, non-convergence, etc.).

In this paper, we avoid the issues of non-convexity by using parameterized penalty functions, which aid in ensuring the strict convexity of the proposed objective function. We propose to solve the following improved sparse low-rank (ISLR) formulation

$$\arg \min_{\mathbf{X} \in \mathbb{R}^{m \times n}} \left\{ F(\mathbf{X}) := \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^k \phi(\sigma_i(\mathbf{X}); a_0) + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n \phi(\mathbf{X}_{i,j}; a_1) \right\}, \quad (3)$$

*Corresponding author. Email address: ankit.parekh@nyu.edu
Source Code available at <https://github.com/aparek/ISLRMatrix>

where $k = \min(m, n)$ and $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is a parameterized non-convex penalty function (see Sec. 2.1). Note that, if $\lambda_1 = 0$, then the ISLR formulation reduces to the generalized nuclear norm minimization problem [36, 40]. Further, if $\lambda_1 = 0$ and $\phi(x; a) = |x|$, then the ISLR problem (3) reduces to the singular value thresholding (SVT) problem.

The contributions of this paper are two-fold. First, we show how to set the parameters a_0 and a_1 to ensure that the function F in (3) is strictly convex. Second, we provide an ADMM based algorithm to solve (3), which utilizes single variable-splitting compared to two variable-splitting as in [53]. We guarantee the convergence of ADMM, provided the scalar augmented Lagrangian parameter μ , satisfies $\mu > 1$.

1.1. Related work

The parameterized non-convex penalty functions used in this paper have designated non-convexity, which enables the overall objective function F in (3) to be strictly convex. In particular, if the parameters a_0 and a_1 exceed their critical value, then the function F in (3) is non-convex. A similar framework of convex objective functions with non-convex regularization was studied for several signal processing applications (see for eg., [50], [18], [33], [42] and the references therein). It was reported that non-convex regularization outperformed convex regularization methods for these applications.

The sparse low-rank (SLR) formulation in (2) is different from the low-rank + sparse decomposition [9], also known as the robust principal component analysis (RPCA). Both the SLR and the RPCA formulations utilize the nuclear norm and the ℓ_1 norm as sparsity-inducing regularizers [56, 55]. The RPCA formulation aims to estimate the matrix, which is the sum of a low-rank and a sparse matrix. Note that, in the case of RPCA, the matrix to be estimated is itself neither sparse or low-rank [12, 11]. In contrast, the SLR problem (2), and the one proposed in this paper, considers the case wherein the matrix to be estimated is simultaneously sparse and low-rank (similar to [24]).

Several well-studied convex optimization algorithms, such as ADMM [25, 1], ISTA/FISTA [2, 22], and proximal gradient methods [17] can be applied to solve convex objective functions of the type (2). The SLR objective function (2), has been solved using Generalized Forward-Backward [44], Incremental Proximal Descent [47] (introduced in [3]), Majorization-Minimization [31], and the Inexact Augmented Lagrangian Multiplier (IALM) method [35]. The IALM method can also be used to solve the SLR problem, although with a different data-fidelity term [53].

2. Preliminaries

We denote vectors and matrices by lower and upper case letters respectively. For a matrix $\mathbf{Y}$, we use the fol-

Figure 1: (a) Non-convex penalty function ϕ in (6) for three values of a . (b) The twice continuously differentiable concave function $s(x; a) = \phi(x; a) - |x|$ in (8) for the corresponding values of a .

lowing entry-wise norms,

$$\|\mathbf{Y}\|_F^2 := \sum_{i,j} |\mathbf{Y}_{i,j}|^2, \quad \|\mathbf{Y}\|_1 := \sum_{i,j} |\mathbf{Y}_{i,j}|. \quad (4)$$

Further, we use the nuclear norm (also called the ‘Schatten-1’ norm) defined as

$$\|\mathbf{Y}\|_* := \sum_{i=1}^k \sigma_i(\mathbf{Y}), \quad (5)$$

where $\sigma_i(\mathbf{Y})$ represent the singular values of the matrix $\mathbf{Y} \in \mathbb{R}^{m \times n}$ and $k = \min(m, n)$.

2.1. Parameterized non-convex penalty functions

We propose to use non-convex penalty functions $\phi(x; a)$ parameterized by the parameter $a \geq 0$. The value of a provides the degree of non-convexity of the penalty functions. Below we define such non-convex penalty functions and list their properties.

Assumption 1. The non-convex penalty function $\phi: \mathbb{R} \rightarrow \mathbb{R}$ satisfies the following

1. ϕ is continuous on $\mathbb{R}$, twice differentiable on $\mathbb{R} \setminus \{0\}$ and symmetric, i.e., $\phi(-x; a) = \phi(x; a)$
2. $\phi'(x) > 0, x > 0$
3. $\phi''(x) \leq 0, x > 0$
4. $\phi'(0^+) = 1$
5. $\inf_{x \neq 0} \phi''(x; a) = \phi''(0^+; a) = -a$

An example of a non-convex penalty function satisfying Assumption 1 is the rational penalty function [23] defined as

$$\phi(x; a) := \frac{|x|}{1 + a|x|/2}, \quad a \geq 0. \quad (6)$$

The ℓ_1 norm is recovered as a special case of the non-convex rational penalty function (i.e., if $a = 0$, then $\phi(x; 0) = |x|$). Figure 1(a) shows the rational penalty function (6) for different values of a . Other examples of penalty functions satisfying Assumption 1 are the logarithmic penalty [10, 38], arctangent penalty [49] and the Laplace penalty [51].

The proximity operator of ϕ [16], $\text{prox}_\phi : \mathbb{R} \rightarrow \mathbb{R}$, is defined as

$$\text{prox}_\phi(y; \lambda, a) := \arg \min_{x \in \mathbb{R}} \left\{ \frac{1}{2}(y - x)^2 + \lambda \phi(x; a) \right\}.$$

The proximity operator associated with the function $\phi(x; a)$, satisfying Assumption 1, is continuous with

$$\text{prox}_\phi(y; \lambda, a) = 0, \forall |y| < \lambda, \quad (7)$$

if $0 \leq a < 1/\lambda$. The proximity operators associated with the arctangent and the logarithmic penalty are provided in [49]. Note that for $a = 0$, the proximity operator is the soft-threshold function [19].

The proximity operator associated with the ℓ_1 norm is the well-known soft-threshold function [19]. Note that the soft-threshold function underestimates non-zero values. In contrast, the proximity operators, associated with the non-convex penalty functions satisfying Assumption 1, approach the identity function asymptotically [49]. Thus, the proximity operators used in this paper estimate the non-zero values more accurately than the ℓ_1 norm.

3. Convexity Condition

In this section we derive a condition to ensure that the function F in (3) is strictly convex. In particular, we show that the objective function F is strictly convex if a_0 and a_1 lie inside a designated region. To this end, we note the following lemmas.

Lemma 1. [41] Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be a non-convex penalty function satisfying Assumption 1. The function $s : \mathbb{R} \rightarrow \mathbb{R}$ defined as

$$s(x; a) := \phi(x; a) - |x|, \quad (8)$$

is twice continuously differentiable, concave and

$$-a \leq s''(x; a) \leq 0. \quad (9)$$

The twice continuously differentiable function $s(x; a) = \phi(x; a) - |x|$ is shown in Fig. 1(b), for three values of a .

Lemma 2. [40] Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be a non-convex penalty function satisfying Assumption 1 and $s : \mathbb{R} \rightarrow \mathbb{R}$ be the function as defined in Lemma 1. The function $G_1 : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ defined as

$$G_1(\mathbf{X}) := \frac{\alpha_1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^k s(\sigma_i(\mathbf{X}); a_0), \quad (10)$$

where $k = \min(m, n)$ and $\alpha_1 > 0$, is strictly convex if

$$0 \leq a_0 < \frac{\alpha_1}{\lambda_0}. \quad (11)$$

Note that the proof of Lemma 2 in [40] considers the case if $\alpha_1 = 1$, however the generalization for $\alpha_1 > 0$ follows directly.

Lemma 3. Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be a non-convex penalty function satisfying Assumption 1 and $s : \mathbb{R} \rightarrow \mathbb{R}$ be the function as defined in Lemma 1. The function $G_2 : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ defined as

$$G_2(\mathbf{X}) := \frac{\alpha_2}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n s(\mathbf{X}_{ij}; a_1), \quad (12)$$

where $\alpha_2 > 0$, is strictly convex if

$$0 \leq a_1 < \frac{\alpha_2}{\lambda_1}. \quad (13)$$

Proof 1. The Frobenius norm of a matrix $\mathbf{Y} \in \mathbb{R}^{m \times n}$ can be viewed as the ℓ_2 norm of a vector $\mathbf{y} \in \mathbb{R}^{mn}$, where $\mathbf{y}$ contains the entries of the matrix $\mathbf{Y}$. Similarly, we stack the entries of the matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ into a vector $\mathbf{x} \in \mathbb{R}^{mn}$ and re-write the function G_2 as

$$G_2(\mathbf{x}) = \frac{\alpha_2}{2} \|\mathbf{y} - \mathbf{x}\|_2^2 + \lambda_1 \sum_{i=1}^{mn} s(x_i, a_1). \quad (14)$$

In order to ensure the strict convexity of G_2 , we seek to ensure that the Hessian of G_2 be positive definite (i.e., $\nabla^2 G \succ 0$). To this end, the Hessian of G_2 is given by

$$\nabla^2 G_2 = \alpha_2 \mathbf{I} + \lambda_1 \cdot \text{diag}(s''(x_1; a_1), \dots, s''(x_{mn}; a_1)), \quad (15)$$

where $\text{diag}(\cdot)$ represents a diagonal matrix. Note that the identity matrix in (15) is of size $mn \times mn$. To ensure that $\nabla^2 G_2$ is positive definite, we seek to ensure

$$\alpha_2 + \lambda_2 s''(t; a_1) > 0, \quad \forall t \in \mathbb{R}, \quad (16)$$

$$s''(t; a_1) > -\frac{\alpha_2}{\lambda_1}. \quad (17)$$

Thus, using Lemma 1 and (17), the Hessian of G_2 , i.e., $\nabla^2 G_2$, is positive definite if $0 \leq a_1 < \alpha_2/\lambda_1$.

The following theorem provides the critical values of the parameters a_0 and a_1 to ensure that the function F in (3) is strictly convex.

Theorem 1. Let $\phi : \mathbb{R} \rightarrow \mathbb{R}$ be a parameterized non-convex penalty function satisfying Assumption 1. The function

$F: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ defined as

$$F(\mathbf{X}) := \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^m \phi(\sigma_i(\mathbf{X}); a_0) + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n \phi(\mathbf{X}_{i,j}; a_1), \quad (18)$$

is strictly convex if

$$0 \leq a_0 \lambda_0 + a_1 \lambda_1 < 1. \quad (19)$$

Proof 2. Let $\alpha \in [0, 1]$. Consider the function $G: \mathbb{R}^{m \times n} \rightarrow \mathbb{R}$ defined as

$$G(\mathbf{X}) := \frac{\alpha}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^m s(\sigma_i(\mathbf{X}); a_0) + \frac{1-\alpha}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n s(\mathbf{X}_{i,j}; a_1), \quad (20)$$

where $s: \mathbb{R} \rightarrow \mathbb{R}$ is defined in Lemma 1. Using the functions G_1 and G_2 , as defined in (10) and (12) respectively, the function G can be written as

$$G(\mathbf{X}) = G_1(\mathbf{X}) + G_2(\mathbf{X}), \quad (21)$$

where $\alpha_1 = \alpha$, and $\alpha_2 = 1 - \alpha$. Due to Lemma 2 and Lemma 3, the function G is strictly convex (being a sum of two strictly convex functions) if a_0 and a_1 satisfy the following inequalities,

$$0 \leq a_0 < \frac{\alpha}{\lambda_0}, \quad (22)$$

$$0 \leq a_1 < \frac{1-\alpha}{\lambda_1}. \quad (23)$$

Combining (22) and (23), we obtain,

$$0 \leq a_0 \lambda_0 + a_1 \lambda_1 < 1. \quad (24)$$

As a result, the function G is strictly convex if a_0 and a_1 satisfy the inequality in (24). Recall that $\phi(x; a) = s(x; a) + |x|$ from (8), due to which the function F in (3) can be written as

$$F(\mathbf{X}) = \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^m [s(\sigma_i(\mathbf{X}); a_0) + |\sigma_i(\mathbf{X})|] + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n [s(\mathbf{X}_{i,j}; a_1) + |\mathbf{X}_{i,j}|] \quad (25)$$

$$= \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^m s(\sigma_i(\mathbf{X}); a_0) + \|\mathbf{X}\|_* + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n s(\mathbf{X}_{i,j}; a_1) + \|\mathbf{X}\|_1 \quad (26)$$

$$= G(\mathbf{X}) + \|\mathbf{X}\|_* + \|\mathbf{X}\|_1. \quad (27)$$

Figure 2: Illustration of the convexity condition provided by Theorem 1. (a) The function G is convex when a_0 and a_1 satisfy the inequality (19). (b) The function G is non-convex otherwise (multiple local minima can be seen in the contour plot).

Figure 3: The function F in (3) is strictly convex for all values of a_0 and a_1 inside the triangular region.

Thus, the function F is strictly convex (being a sum of a strictly convex function and convex functions) if a_0 and a_1 satisfy the inequality (24).

The convexity condition provided by Theorem 1 is illustrated in Fig. 2. The matrix $\mathbf{X}$ is constructed by tiling 10 copies of the matrix $\mathbf{Z} \in \mathbb{R}^{2 \times 2}$,

$$\mathbf{Z} = \begin{bmatrix} x_1 & x_2 \\ x_2 & x_1 \end{bmatrix}, \quad x_i \in \mathbb{R}, \quad (28)$$

and randomly setting 70% of its entries zero. Thus, the matrix $\mathbf{X}$ is of rank 2. We use $\lambda_0 = \lambda_1 = 1$ and set $\mathbf{Y} = \mathbf{0}$. We set the value of $a_0 = 0.8$, and $a_1 = 0.19$, as per (19), to ensure that the function G in (20) is strictly convex. As seen in Fig. 2(a), the function G is strictly convex. However, on increasing the value of a_1 to 1, the function G is non-convex, as seen in Fig. 2(b).

The inequality (19) given by Theorem 1, constitutes a convexity triangle for the function F in (3). For all values of a_0 and a_1 inside the triangular region of Fig. 3 the function F is strictly convex. However, the function

Algorithm 1 Solution to the proposed ISLR problem. The objective function F is given in (3).

```

1: Input:  $\mathbf{Y}$ ,  $\lambda_i$ ,  $a_i$ ,  $\mu$ ,  $\epsilon$ 
2: Initialize:  $\mathbf{Z} = 0$ ,  $\mathbf{D} = 0$ 
3: repeat
4:    $\mathbf{X} \leftarrow \text{prox}_\phi \left( \frac{1}{1+\mu} (\mathbf{Y} + \mu(\mathbf{Z} + \mathbf{D})); \frac{\lambda_1}{1+\mu}, a_1 \right)$ 
5:    $[\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}] \leftarrow \text{SVD}(\mathbf{X} - \mathbf{D})$ 
6:    $\mathbf{Z} \leftarrow \mathbf{U} \cdot \text{prox}_\phi(\mathbf{\Sigma}; \lambda_0/\mu, a_0) \cdot \mathbf{V}^T$ 
7:    $\mathbf{D} \leftarrow \mathbf{D} - (\mathbf{X} - \mathbf{Z})$ 
8: until  $\|F(\mathbf{X}^k) - F(\mathbf{X}^{k-1})\|_2 < \epsilon \|F(\mathbf{X}^k)\|_2$ 

```

F in (3) is non-convex for values of a_0 and a_1 outside the triangular region.

4. Algorithm

We use the alternating direction method of multipliers (ADMM) [6] in conjunction with variable splitting to derive an algorithm for the solution to the ISLR problem (3). The convergence of ADMM to the global minimum is guaranteed when the objective function is a sum of two convex functions [20]. The convergence of ADMM for a non-convex optimization problem to a stationary point is guaranteed under certain mild assumptions [29].

The following theorem derives an algorithm for solving the ISLR problem and guarantees its convergence. In particular, the theorem provides a condition on the value of the scalar augmented Lagrangian parameter μ , to ensure that the sub-problems of ADMM are strictly convex.

Theorem 2. Let $\phi: \mathbb{R} \rightarrow \mathbb{R}$ be a non-convex penalty function satisfying Assumption 1. Let a_0 and a_1 satisfy

$$0 \leq a_0 \lambda_0 + a_1 \lambda_1 < 1, \quad (29)$$

for $\lambda_0, \lambda_1 \geq 0$. Further, let μ be the scalar augmented Lagrangian parameter. If $\mu > 1$, then the iterative algorithm in Table 1 converges to the global minimum of the function F in (3).

Proof 3. Without loss of generality, we set $m = n$. We re-write the ISLR objective function (3) using variable splitting [1] as

$$\begin{aligned} \arg \min_{\mathbf{X}} & \left\{ \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \lambda_0 \sum_{i=1}^m \phi(\sigma_i(\mathbf{Z}); a_0) \right. \\ & \left. + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n \phi(\mathbf{X}_{i,j}; a_1) \right\}, \\ \text{s.t. } & \mathbf{X} = \mathbf{Z}. \end{aligned} \quad (30)$$

The minimization of the ISLR objective function in (30) is separable in $\mathbf{X}$ and $\mathbf{U}$. Applying ADMM to (30), yields

the following iterative procedure, where μ is the scalar augmented Lagrangian parameter and $\mathbf{D} \in \mathbb{R}^{m \times n}$ is update variable.

$$\mathbf{X} \leftarrow \arg \min_{\mathbf{X}} \left\{ \frac{1}{2} \|\mathbf{Y} - \mathbf{X}\|_F^2 + \frac{\mu}{2} \|\mathbf{X} - (\mathbf{Z} + \mathbf{D})\|_F^2 + \lambda_1 \sum_{i=1}^m \sum_{j=1}^n \phi(\mathbf{X}_{i,j}; a_1) \right\} \quad (31a)$$

$$\mathbf{Z} \leftarrow \arg \min_{\mathbf{Z}} \left\{ \frac{\mu}{2} \|\mathbf{X} - \mathbf{D} - \mathbf{Z}\|_F^2 + \lambda_0 \sum_{i=1}^m \phi(\sigma_i(\mathbf{Z}); a_0) \right\} \quad (31b)$$

$$\mathbf{D} \leftarrow \mathbf{D} - (\mathbf{X} - \mathbf{Z}) \quad (31c)$$

Combining the quadratic terms and ignoring the constant terms, the sub-problem (31a) can be written as

$$\mathbf{X} \leftarrow \arg \min_{\mathbf{X}} \left\{ \frac{1}{2} \left\| \frac{1}{1+\mu} (\mathbf{Y} + \mu(\mathbf{Z} + \mathbf{D})) - \mathbf{X} \right\|_F^2 + \frac{\lambda_1}{1+\mu} \sum_{i=1}^m \sum_{j=1}^n \phi(\mathbf{X}_{i,j}; a_1) \right\}. \quad (32)$$

Since $a_1 < 1/\lambda_1$, as per the assumption, the sub-problem (32) is strictly convex and its solution can be obtained using the proximal operator associated with ϕ , i.e.,

$$\mathbf{X} \leftarrow \text{prox}_\phi(\tilde{\mathbf{Y}}; \lambda_1/(1+\mu), a_1). \quad (33)$$

The sub-problem (31b) is the generalized nuclear norm minimization problem, whose solution in closed form is provided by Theorem 2 in [40]. Note that the solution is guaranteed to be the global minimum if $a_0 < \mu/\lambda_0$. Hence, using the inequality (19), the sub-problem (31b) is guaranteed to be strictly convex if $\mu > 1$. As a result, using $(\mathbf{X} - \mathbf{D}) = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ as the singular value decomposition (SVD) of the matrix $\mathbf{X} - \mathbf{D}$, we get

$$\mathbf{Z} \leftarrow \mathbf{U} \cdot \text{prox}_\phi(\mathbf{\Sigma}; \lambda_0/\mu, a_0) \cdot \mathbf{V}^T. \quad (34)$$

Combining (33) and (34), we obtain the iterative algorithm 1, which converges to the global minimum of the ISLR objective function in (3).

5. Examples

We illustrate the proposed ISLR method for estimating simultaneously sparse and low-rank matrices via the following examples. We first describe setting the parameters for the proposed method and then showcase the examples.

5.1. Parameter tuning

The proposed ISLR algorithm 1 requires the specification of two regularization parameters λ_0 and λ_1 , two penalty parameters a_0 and a_1 and the scalar augmented

Figure 4: The value of the objective function F in (3) at every iteration of the ISLR algorithm for the speech signal denoising problem.

Figure 5: Log-log plot showing the convergence of the ISLR objective function (3) and the SLR objective function (2).

Lagrangian parameter μ . We set the regularization parameters λ_i as

$$\lambda_i = \beta_i \sigma, \quad i = 0, 1, \quad (35)$$

where σ is the standard deviation of AWGN (or an estimate) and β_i are chosen so as to maximize the signal-to-noise ratio (SNR) for the SLR and the ISLR methods. The values of a_0 and a_1 are set as

$$a_0 = \frac{c}{\lambda_0}, \quad c \in (0, 1) \quad (36)$$

$$a_1 = \frac{1 - a_0 \lambda_0}{\lambda_1}, \quad (37)$$

respectively. The value of $a_0 \in [0, 1/\lambda_0]$ affects the sparsity of the singular values, and the value of $a_1 \in [0, 1/\lambda_1]$ affects the sparsity of the elements of the matrix to be estimated. Thus, if the sparsity of the singular values is favored over the sparsity of the elements of the matrix to be estimated, the point (a_0, a_1) may be set in the lower-right region of the convexity triangle shown in Fig. 3. Alternatively, if the sparsity of the elements is preferred over the sparsity of the singular values of the matrix to be estimated, the point (a_0, a_1) may be set in the upper-left region of convexity triangle in Fig. 3.

We set the value of μ as $\mu = 1.5$. As per Theorem 2, the ISLR algorithm listed in Table 1 is guaranteed to converge

Figure 6: (a) Average RSE as a function of the rank of the input matrix. (b) Average RSE as a function of the level of sparsity of the input matrix.

to the global minimum for $\mu > 1$. However, depending on the value of μ , the convergence may be slow. As such for this example, and the one that follows, we run the ISLR algorithm till a certain tolerance level is reached, i.e., we run the algorithm till

$$\|F(\mathbf{X}^k) - F(\mathbf{X}^{k-1})\|_2 < \epsilon \|F(\mathbf{X}^k)\|_2, \quad (38)$$

where ϵ is a user-defined tolerance level, usually set to $\epsilon = 10^{-5}$. The value of the objective function (3) for 20 iterations of the ISLR algorithm, with different values of μ , is shown in Fig. 4. Fig. 5 shows the log-log plot of the stopping criteria for the ISLR objective function and the SLR objective function. Note that for the examples that follow, we use the arctangent penalty [49] as the nonconvex penalty for the proposed ISLR method.

5.2. Synthetic data

We generate a synthetic matrix $\mathbf{M} \in \mathbb{R}^{m \times n}$ of rank k using two random matrices $\mathbf{A} \in \mathbb{R}^{m \times k}$ and $\mathbf{B} \in \mathbb{R}^{k \times n}$ such that

$$\mathbf{M} := \mathbf{A} \cdot \mathbf{B}, \quad (39)$$

where the entries of $\mathbf{A}$ and $\mathbf{B}$ are chosen from an i.i.d standard normal distribution. To measure the performance of the proposed ISLR method and the SLR method, we use the normalized root square error (RSE) defined as

$$\text{RSE} := \frac{\|\mathbf{X}_{\text{est}} - \mathbf{X}_{\text{org}}\|_F}{\|\mathbf{X}_{\text{org}}\|_F}, \quad (40)$$

where $\mathbf{X}_{\text{est}}$ represents the estimated matrix and $\mathbf{X}_{\text{org}}$ represents the desired clean matrix.

We consider two types of simulations: RSE as a function of the rank (k) of the synthetically generated matrix and RSE as a function of the sparsity level of the input matrix. For the first simulation, we fix the sparsity level at 60% (i.e., approximately 40% of the entries of the clean input matrix $\mathbf{M}$ are zero) while varying the rank k of the input matrix. We add white Gaussian noise ($\sigma = 0.2$) to $\mathbf{M}$ to generate a noisy input matrix $\mathbf{Y}$. We generate 15 matrices for each value of the rank k where $1 \leq k \leq 100$ in

increments of 5 and denoise them using the proposed ISLR method (3) and the SLR method (2). For the second simulation, wherein we consider the RSE as a function of the level of sparsity of the input matrix, we synthetically generate 15 matrices $\mathbf{M}$ of fixed rank $k = 10$ but with varying levels of sparsity (from 10% to 90%). Again, we add white Gaussian noise ($\sigma = 0.2$) to generate the noisy matrix $\mathbf{Y}$.

Figure 6(a) shows the average RSE values as a function of the rank k of the input matrix. Figure 6(b) shows the average RSE values as a function of the level of sparsity of the input matrix. The proposed ISLR method consistently obtains lower RSE values than the SLR method. As expected, for both the methods, the proposed ISLR method and the SLR method, RSE values are lower for matrices that are relatively more sparse. On the other hand, as seen in Fig. 6(a), we observe that for matrices that are not necessarily low-rank but have decaying singular values, lower values of RSE are obtained for both the methods. Note that for both the simulations, we do a grid search over a range of values for β_0 and β_1 to obtain the values of λ_0 and λ_1 respectively, which yield the lowest RSE values (recall that $\lambda_i = \beta_i \sigma$, for $i = 0, 1$). Furthermore, in both the simulations, we fix the value of c at $c = 0.5$ for setting the values of a_0 and a_1 .

5.3. Speech signal denoising

We consider the problem of denoising a speech signal with AWGN. We apply the sparse low-rank matrix estimation methods to the spectrogram of the noisy speech signal, and invert the estimated spectrogram to the time domain to obtain the denoised speech signal. Specifically, if $\mathbf{y}$ is the input speech signal, then the denoised estimate $\mathbf{x}^*$ is obtained using ISLR as

$$\mathbf{x}^* = \mathbf{S}^\dagger \left\{ \text{ISLR}(\mathbf{S}(\mathbf{y}); \lambda_i, a_i) \right\}, \quad i = 0, 1, \quad (41)$$

where $\mathbf{S}$ and $\mathbf{S}^\dagger$ represent the short-time Fourier transform (STFT) and its inverse, respectively. For this example, we set $\mathbf{S}$ to be an over-complete STFT, implemented with perfect reconstruction, i.e., $\mathbf{S}^\dagger \mathbf{S} = \mathbf{I}$. The STFT is implemented with 50% overlap between the windows, for a window size of 64 samples. The DFT length is set to 512 samples. We add AWGN ($\sigma = 0.03$) to realize the noisy speech signal. We compare the proposed ISLR method with the SLR method (2) and the nonconvex sparse low-rank matrix estimation method [24] which uses the weighted nuclear norm [26] and the weighted ℓ_1 norm [10].

The spectrogram of the clean speech signal is shown in Fig. 7(a). It can be seen that the spectrogram consists of repeated ridges. The spectrogram of the noisy speech signal is shown in Fig. 7(b). Figure 7(c) shows the denoised spectrogram obtained using the SLR method (2) and Fig. 7(d) shows the denoised spectrogram using the ISLR method (3). Figure 7(e) shows the spectrogram estimated using the nonconvex method with weighted ℓ_1 norm and the weighted nuclear norm (see Algorithm 1 in

Figure 7: Illustration of proposed method for denoising an audio signal. The colorbar values are in dB.

Figure 8: SNR as a function of λ_0 with fixed λ_1 for the ISLR and the SLR methods.

[24]). The ISLR estimated spectrogram has a higher SNR than the SLR estimated spectrogram and contains fewer artifacts. The nonconvex method obtains a slightly higher SNR than the proposed ISLR method and also contains relatively fewer artifacts than the SLR method. The proposed ISLR method obtains SNR values comparable to the state-of-the-art nonconvex method, while being able to guarantee convergence to the unique global minimum.

Figure 9: Illustration of denoising the weighted adjacency matrix representing protein interactions in the ‘Escherichia coli’ bacteria. Note that ‘nz’ represents the number of non-zero elements in the matrix.

Figure 8 shows the SNR as a function of the regularization parameter λ_1 , when λ_0 is fixed, for the SLR and the ISLR methods. Note that the standard deviation σ of the noise level is also fixed at ($\sigma = 0.03$). The improvement in the SNR value, when using the ISLR method, is the same when the value of λ_0 is also varied, in addition to the value of λ_1 . The SNR values are obtained by averaging over 15 realizations for each λ_0, λ_1 pair.

5.4. Protein Interactions

Protein interactions in the ‘Escherichia coli’ bacteria, scored by strength in $[0, 2]$, were studied in [30]. The data can be represented as a weighted graph, which is sparse and low-rank by nature [47]. The rationale behind the low-rank property of the weighted graph, is that the interactions between two sets of proteins are governed by a small set of factors [47, 5].

Figure 9(a) shows the protein interaction data as a weighted adjacency matrix. The adjacency matrix is obtained after retaining 440 proteins of the entire set of 4394 proteins. We corrupt 10% of the entries of the clean adjacency matrix, selected uniformly at random, with uniform noise in the interval $[0, \sigma]$. The noisy adjacency matrix is shown in Fig. 9(b), with $\sigma = 0.3$. We set the parameters λ_i and $a_i, i = 1, 2$, as in the previous example. Figure 9(c) and Fig. 9(d) show the denoised adjacency matrices obtained using the ISLR and the SLR methods, respectively. As in the case of previous example, the ISLR method offers a better RSE, and tends to correctly estimate the sparsity-pattern of the true matrix.

Figure 10: Average RSE as a function of σ .

In order to assess the relative performance of the proposed ISLR method (3) in comparison to the SLR method (2), we realize 15 noisy adjacency matrices and denoise them. For each value of σ , we choose the parameters λ_i , for both the methods, that yields the lowest RSE. Shown in Fig. 10 are the average RSE values as a function of σ . It can be seen that the ISLR method consistently offers a lower RSE.

6. Conclusion

We consider the problem of estimating a sparse low-rank matrix from its noisy observation. We generalize the convex formulation proposed for estimation of sparse low-rank matrix estimation [47], by utilizing non-convex sparsity-inducing regularizers. The non-convex penalty functions proposed are known to estimate the non-zero signal values more accurately. We show how to set the parameters of the non-convex penalty functions, so as to preserve the convexity of the overall problem (sum of data-fidelity and the regularization terms). The critical value of the non-convex penalty parameters define a convexity triangle; for all values of the nonconvex penalty parameters within this triangular region, the cost function is guaranteed to be strictly convex.

We derive an efficient algorithm using ADMM with a single variable-splitting which solves the proposed convex objective function consisting of non-convex regularizers. We guarantee the convergence of ADMM to the global minimum of the objective function, provided the scalar augmented Lagrangian parameter μ is chosen such that $\mu > 1$. We illustrate several examples to demonstrate the effectiveness of the proposed formulation for estimation of sparse low-rank matrices.

The proposed method utilizes separable penalty functions (nonconvex) to induce sparsity stronger than separable convex penalty functions. A possible future direction involves the use of non-separable penalty functions, possibly nonconvex, that are designed so as to ensure the strict convexity of the objective function [48].

7. Acknowledgements

The authors thank the anonymous reviewers for their detailed suggestions and corrections. This work was sup-

ported by the ONR under grant N00014-15-1-2314 and the NSF under grant CCF-1525398.

8. References

- [1] M. Afonso, J. Bioucas-Dias, and M. Figueiredo. Fast image recovery using variable splitting and constrained optimization. *IEEE Trans. Image Process.*, 19(9):2345–2356, sep 2010.
- [2] A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sci.*, 2(1):183–202, 2009.
- [3] D. P. Bertsekas. Incremental Gradient, subgradient, and proximal methods for convex optimization :A survey. *Optimization*, 2010(December):1–38, 2010.
- [4] J. Bien and R. J. Tibshirani. Sparse estimation of a covariance matrix. *Biometrika*, 98(4):807–820, 2011.
- [5] J. R. Bock and D. A. Gough. Predicting protein–protein interactions from primary structure. *Bioinformatics*, 17(5):455–460, 2001.
- [6] S. Boyd, N. Parikh, E. Chu, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, 3(1):1–122, 2010.
- [7] A. Buja, Z. Ma, and D. Yang. Optimal denoising of simultaneously sparse and low rank matrices in high dimensions. *Proc. Allert. Conf. Commun. Control Comput.*, pages 445–447, oct 2013.
- [8] F. Bunea, Y. She, and M. H. Wegkamp. Joint variable and rank selection for parsimonious estimation of high-dimensional matrices. *Ann. Stat.*, 40(5):2359–2388, 2012.
- [9] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):1–37, 2011.
- [10] E. J. Candès, M. B. Wakin, and S. P. Boyd. Enhancing sparsity by reweighted l1 minimization. *J. Fourier Anal. Appl.*, 14(5):877–905, dec 2008.
- [11] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-Sparsity Incoherence for Matrix Decomposition. *SIAM J. Optim.*, 21(2):572–596, 2009.
- [12] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Sparse and low-rank matrix decompositions. In *2009 47th Annu. Allert. Conf. Commun. Control Comput.*, pages 962–967, sep 2009.
- [13] R. Chartrand. Fast algorithms for nonconvex compressive sensing: MRI reconstruction from very few data. In *IEEE Int. Symp. Biomed. Imag.*, pages 262–265, jul 2009.
- [14] R. Chartrand. Nonconvex splitting for regularized low-rank + sparse decomposition. *IEEE Trans. Signal Process.*, 60(11):5810–5819, 2012.
- [15] L. Chen and J. Z. Huang. Sparse Reduced-Rank Regression for Simultaneous Dimension Reduction and Variable Selection. *J. Am. Stat. Assoc.*, 107(500):1533–1545, 2012.
- [16] P. L. Combettes and J.-C. Pesquet. Proximal thresholding algorithm for minimization over orthonormal bases. *SIAM J. Optim.*, 18(4):1351–1376, nov 2007.
- [17] P. L. Combettes and J.-C. Pesquet. Proximal splitting methods in signal processing. In H. H. Bauschke, editor, *Fixed-Point Algorithms Inverse Probl. Sci. Eng.*, pages 185–212. Springer-Verlag, 2011.
- [18] Y. Ding and I. W. Selesnick. Artifact-free wavelet denoising: non-convex sparse regularization, convex optimization. *IEEE Signal Process. Lett.*, 22(9):1364–1368, sep 2015.
- [19] D. L. Donoho. De-noising by soft-thresholding. *IEEE Trans. Inf. Theory*, 41(3):613–627, 1995.
- [20] J. Eckstein and D. P. Bertsekas. On the Douglas-Rachford splitting method and the proximal point algorithm for maximal monotone operators. *Math. Program.*, 55(1-3):293–318, apr 1992.
- [21] N. El Karoui. Operator norm consistent estimation of large-dimensional sparse covariance matrices. *Ann. Stat.*, 36(6):2717–2756, 2008.
- [22] M. A. T. Figueiredo, J. M. Bioucas-Dias, and R. D. Nowak. Majorization-minimization algorithms for wavelet-based image restoration. *IEEE Trans. Image Process.*, 16(12):2980–2991, dec 2007.
- [23] D. Geman and G. Reynolds. Constrained restoration and the recovery of discontinuities. *IEEE Trans. Pattern Anal. Mach. Intell.*, 14(3):367–383, mar 1992.
- [24] P. V. Giampouras, K. E. Themelis, A. A. Rontogiannis, and K. D. Koutroumbas. Simultaneously Sparse and Low-Rank Abundance Matrix Estimation for Hyperspectral Image Unmixing. *IEEE Trans. Geosci. Remote Sens.*, 54(8):4775–4789, aug 2016.
- [25] T. Goldstein and S. Osher. The Split Bregman Method for L1-Regularized Problems. *SIAM J. Imaging Sci.*, 2(2):323–343, jan 2009.
- [26] S. Gu, L. Zhang, W. Zuo, and X. Feng. Weighted nuclear norm minimization with application to image denoising. In *IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 2862–2869, jun 2014.
- [27] L. Han and X.-L. Liu. Convex Relaxation Algorithm for a Structured Simultaneous Low-Rank and Sparse Recovery Problem. *J. Oper. Res. Soc. China*, 3(3):363–379, 2015.
- [28] A. Hansson, Z. Liu, and L. Vandenberghe. Subspace system identification via weighted nuclear norm optimization. *Proc. IEEE Conf. Decis. Control*, pages 3439–3444, 2012.
- [29] M. Hong, Z.-Q. Lo, and M. Razaviyayn. Convergence analysis of alternating direction method of multipliers for a family of nonconvex problems. *Proc. IEEE Int. Conf. Acoust. Speech Signal Process. ICASSP*, pages 1–5, 2015.
- [30] P. Hu, S. C. Janga, M. Babu, J. J. Díaz-Mejía, and G. Butland. Global functional atlas of Escherichia coli encompassing previously uncharacterized proteins. *PLoS Biol.*, 7(4):0929–0947, 2009.
- [31] Y. Hu, S. Lingala, and M. Jacob. A fast majorizeminimize algorithm for the recovery of sparse and low-rank matrices. *IEEE Trans. Image Process.*, 21(2):742–753, 2012.
- [32] V. Jojic, S. Saria, and D. Koller. Convex envelopes of complexity controlling penalties: the case against premature envelopment. *Proc. Conf. Artif. Intell. Stat. AISTATS*, 15:399–406, 2011.
- [33] A. Lanza, S. Morigi, and F. Sgallari. Convex image denoising via non-convex regularization. In J.-F. Aujol, M. Nikolova, and N. Papadakis, editors, *Scale Sp. Var. Methods Comput. Vis.*, volume 9087 of *Lecture Notes in Computer Science*, pages 666–677. Springer, 2015.
- [34] M. Lee, H. Shen, J. Z. Huang, and J. S. Marron. Biclustering via Sparse Singular Value Decomposition. *Biometrics*, 66(4):1087–1095, 2010.
- [35] Z. Lin, M. Chen, and Y. Ma. The Augmented Lagrange Multiplier Method for Exact Recovery of Corrupted Low-Rank Matrices. *arXiv:1009.5055*, pages 1–23, 2010.
- [36] C. Lu, J. Tang, S. Yan, and Z. Lin. Generalized nonconvex nonsmooth low-rank minimization. In *IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 4130–4137, jun 2014.
- [37] C. Lu, C. Zhu, C. Xu, S. Yan, and Z. Lin. Generalized singular value thresholding. *arXiv1412.2231 Prepr.*, 2014.
- [38] K. Mohan and M. Fazel. Iterative Reweighted Algorithms for Matrix Rank Minimization. *J. Mach. Learn. Res.*, 13:34413473, 2012.
- [39] R. R. Nadakuditi. OptShrink: An algorithm for improved low-rank signal matrix denoising by optimal, data-driven singular value shrinkage. *IEEE Trans. Inf. Theory*, 60(5):3002–3018, 2014.
- [40] A. Parekh and I. Selesnick. Enhanced low-rank matrix approximation. *IEEE Signal Process. Lett.*, 23(4):493–497, 2016.
- [41] A. Parekh and I. W. Selesnick. Convex denoising using non-convex tight frame regularization. *IEEE Signal Process. Lett.*, 22(10):1786–1790, 2015.
- [42] A. Parekh and I. W. Selesnick. Convex Fused Lasso Denoising with Non-Convex Regularization and its use for Pulse Detection. In *IEEE Symp. Signal Process. Med. Biol.*, pages 1–6,

2015.

- [43] J. Portilla. Image restoration through L0 analysis-based sparse optimization in tight frames. *Proc. IEEE Int. Conf. Image Process. ICIP*, pages 3909–3912, nov 2009.
- [44] H. Raguét, J. Fadili, and G. Peyré. A generalized forward-backward splitting. *SIAM J. Imaging Sci.*, (2):1–29, 2013.
- [45] A. Repetti, E. Chouzenoux, and J.-C. Pesquet. A nonconvex regularized approach for phase retrieval. *Proc. IEEE Int. Conf. Image Process. ICIP*, pages 1753–1757, oct 2014.
- [46] E. Richard, S. Gaïffas, and N. Vayatis. Link Prediction in Graphs with Autoregressive Features. *J. Mach. Learn. Res.*, 15(1):565–593, jan 2014.
- [47] E. Richard, E. C. Paris, N. Vayatis, and P.-A. Savalle. Estimation of Simultaneously Sparse and Low Rank Matrices. *Proc. Int. Conf. Mach. Learn.*, pages 1351–1358, 2012.
- [48] I. Selesnick. Total Variation Denoising Via the Moreau Envelope. *IEEE Signal Process. Lett.*, 24(2):216–220, feb 2017.
- [49] I. Selesnick and I. Bayram. Sparse signal estimation by maximally sparse convex optimization. *IEEE Trans. Signal Process.*, 62(5):1078–1092, mar 2014.
- [50] I. Selesnick, A. Parekh, and I. Bayram. Convex 1-D total variation denoising with non-convex regularization. *IEEE Signal Process. Lett.*, 22(2):141–144, feb 2015.
- [51] J. Trzasko and A. Manduca. Highly undersampled magnetic resonance image reconstruction via homotopic $l(0)$ -minimization. *IEEE Trans. Med. Imaging*, 28(1):106–121, 2009.
- [52] B. Xin, Y. Tian, Y. Wang, and W. Gao. Background subtraction via generalized fused lasso foreground modeling. *arXiv Prepr. arXiv1504.03707*, pages 1–9, 2014.
- [53] T. Zhang, B. Ghanem, S. Liu, C. Xu, and N. Ahuja. Low-Rank Sparse Coding for Image Classification. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 281–288, 2013.
- [54] S.-L. Zhou, N.-H. Xiu, Z.-Y. Luo, and L.-C. Kong. Sparse and Low-Rank Covariance Matrix Estimation. *J. Oper. Res. Soc. China*, (2015):231–250, 2014.
- [55] T. Zhou and D. Tao. GoDec : Randomized low-rank & sparse matrix decomposition in noisy case. *Proc. Int. Conf. Mach. Learn.*, pages 1–8, 2011.
- [56] X. Zhou, C. Yang, H. Zhao, and W. Yu. Low-Rank Modeling and Its Applications in Image Analysis. *ACM Comput. Surv.*, 47(2):36:1–36:33, 2014.