

# Relaxations for Non-Separable Cardinality/Rank Penalties

Carl Olsson<sup>1,2</sup> Daniele Gerosa<sup>1</sup> Marcus Carlsson<sup>1</sup>

<sup>1</sup>Centre for Mathematical Sciences, <sup>2</sup>Department of Electrical Engineering,  
Lund University Chalmers University of Technology

## Abstract

*Rank and cardinality penalties are hard to handle in optimization frameworks due to non-convexity and discontinuity. Strong approximations have been a subject of intense study and numerous formulations have been proposed. Most of these can be described as separable, meaning that they apply a penalty to each element (or singular value) based on size, without considering the joint distribution. In this paper we present a class of non-separable penalties and give a recipe for computing strong relaxations suitable for optimization. In our analysis of this formulation we first give conditions that ensure that the globally optimal solution of the relaxation is the same as that of the original (unrelaxed) objective. We then show how a stationary point can be guaranteed to be unique under the restricted isometry property (RIP) assumption.*<sup>1</sup>

## 1. Introduction

Sparsity and low rank priors are common ways of regularizing ill-posed inverse problems. In the computer vision community they have been employed in a wide variety of applications such as outlier detection/removal, face recognition, rigid and non rigid structure from motion, photometric stereo and optical flow [46, 38, 43, 5, 47, 19, 2, 20]. The prior is typically formulated either as a soft penalty, resulting in a trade-off between data fit and regularization, or as a hard constraint enforcing a particular cardinality/rank. In this paper we formulate the general sparsity regularized problem as

$$G(\text{card}(\mathbf{x})) + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2, \quad (1)$$

<sup>1</sup>This work has been supported by the Swedish Research Council (grant 2018-05375), the Swedish Foundation for Strategic Research (Semantic Mapping and Visual Navigation for Smart Robots), the Wallenberg AI, Autonomous Systems and Software Program (WASP), eSSSENCE and the Crafoord Foundation.

where the function  $G$  can be written as

$$G(k) = \sum_{i=0}^k g_i, \quad (2)$$

where  $0 = g_0 \leq g_1 \leq g_2 \leq \dots \leq g_n \leq \infty$ . Note that  $g_i = \infty$  is allowed (if  $i > 0$ ). This formulation covers both hard and soft priors. If we for example chose  $G(k) = \mu k$  we get an objective that penalizes but does not restrict the sparsity of the solution. In contrast, if we let  $G(k) = 0$  when  $k \leq r$  and  $\infty$  when  $k > r$  we get a hard cardinality constraint. Many other choices for  $G$  are possible. In this paper we will also consider the corresponding matrix version of (1), formulated as

$$G(\text{rank}(X)) + \|\mathbf{A}X - \mathbf{b}\|^2, \quad (3)$$

where  $\mathcal{A}$  is a linear operator. The theory for the vector and sparsity formulations are with a few exceptions very similar, since the rank of a matrix is basically a sparsity prior on the singular values of the matrix. We will therefore state our main results in the vector setting but emphasize that they apply for the matrix setting as well.

In general  $G$  is convex and non-decreasing on the non-negative integers. However as function of the unknown  $\mathbf{x}$ ,  $G(\text{card}(\mathbf{x}))$  is highly non-convex as well as discontinuous and in general these problems are NP-hard [34, 21]. Therefore relaxations have to be employed. In recent years there have been a lot of work on convex as well as non-convex relaxations for both sparsity and rank regularized problems. The standard method is to replace  $\text{card}(\mathbf{x})$  with the convex  $\ell_1$  norm  $\|\mathbf{x}\|_1$  [45, 44, 8, 15]. Furthermore, if the RIP constraint

$$(1 - \delta_r)\|\mathbf{x}\|^2 \leq \|\mathbf{A}\mathbf{x}\|^2 \leq (1 + \delta_r)\|\mathbf{x}\|^2, \quad (4)$$

holds for all vectors  $\mathbf{x}$  with  $\text{card}(\mathbf{x}) \leq r$ , asymptotic performance guarantees can be derived [8]. On the other hand the  $\ell_1$  approach suffers from a shrinking bias since it penalizes both small elements of  $\mathbf{x}$ , assumed to stem from measurement noise, and large elements, assumed to make up the true signal, equally. Hence the suppression of noise also requires an equal

suppression of signal [17, 31]. This insight has led to a large number of non-convex alternatives able to penalize small components proportionally harder than the large ones e.g. [17, 31, 4, 3, 13, 41, 50, 28, 49, 48, 30]. With a few exceptions e.g. [28, 29] global optimality guarantees are generally not available for these formulations. In addition these typically employ separable formulations, that is, a non-convex penalty is applied to each element without regarding the joint element values. Such a formulation can for example not add hard thresholds on the number of non-zero elements of the vector. It is important to note that under RIP the matrix  $A$  typically has a nullspace containing dense vectors. Under such conditions, separable formulations that don't have shrinking bias, often have local minimizers, see Section 2. Hence a non-separable  $G$ , able to strongly penalize high cardinality solutions, is likely to provide better relaxations. On the other hand these are harder to analyse and less common in the literature. The  $k$ -support norm studied in [1, 32] is a non-separable surrogate for the rank function. It is however a convex norm and therefore suffers from a shrinking bias similar to the  $\ell_1$  norm.

The theory of rank minimization is largely analogous that of sparsity. In this context the rank function is typically replaced with the convex nuclear norm  $\|X\|_* = \sum_i \sigma_i(X)$  [42, 6]. In [42] the notion of RIP was generalized to the matrix setting. A number of generalizations that give performance guarantees for the matrix case have appeared [40, 6, 7] and non-convex alternatives have also been considered [35, 39, 33, 25]. The analogue of the  $k$ -support norm was considered in [16, 22]. The so called weighted nuclear norm is a popular choice for vision problems [24, 23, 26]. We are however not aware of any global recovery guarantees with this regularizer. (Note that even though the weighted nuclear norm is linear in the singular value vector it is not a separable nor convex penalty since the singular values are non-linear functions of the matrix elements.) In this work we study the class of non-convex non-separable relaxations of the objective described in (1) and (3). The relaxation is obtained when replacing the regularizer with its quadratic envelope [9]. If  $f(x) = G(\text{card}(\mathbf{x})) + \|\mathbf{x}\|^2$  then the quadratic envelope  $\mathcal{R}_g$  can be defined by

$$\mathcal{R}_g(\mathbf{x}) = f^{**}(\mathbf{x}) - \|\mathbf{x}\|^2, \quad (5)$$

where  $f^{**}(\mathbf{x})$  is the convex envelope of  $f(\mathbf{x})$ . Thus we first add a quadratic penalty to the regularizer, then take the convex envelope and subtract the quadratic function. The intuition behind the choice of regularizer is that  $\mathcal{R}_g(\mathbf{x}) + \|\mathbf{x} - \mathbf{b}\|^2$  is the convex envelope of  $G(\text{card}(\mathbf{x})) + \|\mathbf{x} - \mathbf{b}\|^2$ , see [27], and therefore any

stationary point is a global minimizer. Under RIP the term  $\|\mathbf{x} - \mathbf{b}\|^2$  is likely to behave similarly to  $\|A\mathbf{x} - \mathbf{b}\|^2$  for vectors with  $\text{card}(\mathbf{x}) < r$ . In this paper we formally study the properties of stationary points of the resulting minimizers and give conditions that guarantee global optimality of a stationary point. Note that our work exclusively deals with properties of the objective function and do not assume any particular optimization method. Any method that reaches a stationary point or a local minimum will suffice. The theory presented in this paper unifies and makes significant extensions of the results in [36, 37] where two special cases of the framework are studied.

## 1.1. Relaxations

In this section we give a very brief presentation of the regularizer that we use, which is taken from [27]. The function  $G(\text{card}(\mathbf{x}))$  only depends on the sorted magnitudes of the elements in  $\mathbf{x}$ , and this is also the case for  $\mathcal{R}_g(\mathbf{x})$ . We will denote these by  $\tilde{\mathbf{x}}$  and assume that  $\mathbf{x} = D_s \pi \tilde{\mathbf{x}}$ . Here the vector  $\mathbf{s}$  contains elements that are either  $-1$  or  $1$ ,  $D_s$  is a diagonal matrix with the elements of  $\mathbf{s}$  on the diagonal and  $\pi$  is a permutation matrix. With this notation our regularizer can be written

$$\mathcal{R}_g(\mathbf{x}) = \max_{\tilde{\mathbf{z}}} 2\langle \tilde{\mathbf{x}}, \tilde{\mathbf{z}} \rangle - \sum_{i=1}^n \max(\tilde{z}_i^2 - g_i, 0) - \|\tilde{\mathbf{x}}\|^2, \quad (6)$$

see (42) in the supplementary material for further details. Evaluation of  $\mathcal{R}_g$  requires solving the maximization over  $\tilde{\mathbf{z}}$ . This can be done very fast (logarithmic time in the number of elements of  $\mathbf{x}$ ). A simple (linear time) algorithm is presented in [27].

The relaxation of (1) can then be written

$$\mathcal{R}_g(\tilde{\mathbf{x}}) + \|A\mathbf{x} - \mathbf{b}\|^2. \quad (7)$$

The theory for the matrix case is largely identical to that of the vector case. Here the regularizer only depends on the (sorted) singular values of  $X$  which we will also denote  $\tilde{\mathbf{x}}$ . The relationship between  $X$  and  $\tilde{\mathbf{x}}$  is now the singular value decomposition (SVD)  $X = U D_{\tilde{\mathbf{x}}} V^T$ , where  $U$  and  $V$  are orthogonal matrices. We will therefore write the relaxation of (3)

$$\mathcal{R}_g(\tilde{\mathbf{x}}) + \|AX - \mathbf{b}\|^2. \quad (8)$$

Note that the orthogonal matrices  $U$  and  $V$  are not unique if  $\tilde{\mathbf{x}}$  has elements that are zero. In the vector case we have a similar non-uniqueness in the matrix  $D_s \pi$ .

## 2. Motivation

Being able to use a general prior  $G$  has the benefit that we can design accurate formulations for finding the correct cardinality. While separable regularization considers the size of each variable separately the ability to apply a non-separable prior makes it possible to heavily penalize unlikely solutions. The extreme example, the so called fixed cardinality penalty

$$G(k) = \begin{cases} 0 & k \leq k_{\max} \\ \infty & k > k_{\max} \end{cases}, \quad (9)$$

rules out solutions with more non-zero elements than  $k_{\max}$ , which cannot be achieved with a separable formulation. Less restrictive variants that regard high cardinality states unlikely (but not impossible) can also be used.

The use of a non-separable prior is not only important for modelling purposes but also effects optimization algorithms since separable formulations more often get stuck in local minima. To understand this consider the simple one dimensional problem  $|x|_0 + (x - y)^2$ , where  $|x|_0 = 0$  if  $x = 0$  and 1 otherwise. The goal of this formulation is recover  $x = y$  if  $y$  is large enough not to be considered noise. By taking derivatives it is easy to see that the solution to this problem is given by  $x = y$  if  $|y| \geq 1$  and  $x = 0$  otherwise. Now suppose that  $|x|_0$  is replaced by a function  $f(|x|)$ . It is easy to see that the minimizer of  $f(|x|) + (x - y)^2$  is either  $x = 0$  or a stationary point fulfilling  $x = y - \text{sign}(y) \frac{f'(|x|)}{2}$ . Hence to really recover  $x = y$  when  $y$  is large enough we have to have  $f'(|x|) = 0$ , that is,  $f$  has to be constant for large values. If this is not the case  $f$  will favour smaller solutions resulting in a shrinking bias.

On the other hand a separable regularizer of the form  $F(\mathbf{x}) = \sum_i f(x_i)$ , where  $f$  is constant for large values is likely to have local minima even under RIP. A 2D example depicting the shape of the regularizer is shown in Figure 1. To the left is  $\mathcal{R}_g$  with  $g_1 = g_2 = 1$  which is separable and obtained by applying the function to the right in Figure 1 to both coordinates. For general dimensions these types of penalties yield regularizers that are constant for dense vectors that are large enough. Suppose now that  $\mathbf{x}^*$  minimizes  $\|A\mathbf{x} + \mathbf{b}\|^2$ , that is  $\mathbf{x}^*$  is a least squares solution, and there is a dense vector  $\mathbf{y}$  in the nullspace of  $A$ . Then  $\mathbf{x}^* + \lambda\mathbf{y}$  is also a least squares solution. If we make  $\lambda$  large enough the elements of the vector  $\mathbf{x}^* + \lambda\mathbf{y}$  will be located in the region where the separable relaxation is constant while at the same time minimizing the data term and therefore it is a local minimizer of the relaxation (as well as the original unrelaxed formulation). In [36] conditions that guarantee uniqueness of sparse stationary points under the regularization  $\mathcal{R}_g$

and  $g_1 = g_2 = \dots = g_n$  when RIP holds were given. However, dense stationary points could not be ruled out. In our experimental section we confirm that optimization starting from a least squares solution often results in convergence to poor dense solutions, see Section 4. One way to address this problem is to accept a modest shrinking bias to make sure that the gradient of the regularizer does not vanish for large elements [10]. For the type of local minima described above this is likely to solve the problem, however it is not clear if there are other types of dense minima as well. In this work we instead consider non-separable formulations where dense vectors can be penalized harder. In the middle of Figure 1 we show  $\mathcal{R}_g$  with  $g_1 = 1$  and  $g_2 = \infty$ . In the latter case the  $G$  function excludes dense vectors. This is reflected in the shape of the relaxation  $\mathcal{R}_g$ , which will clearly try to discourage cardinality 2. For vectors of cardinality 1 (and 0) the two options (Left and Middle) provide identical penalties. Our main results give conditions that are sufficient for guaranteeing uniqueness of stationary points among both sparse and dense vectors.




Figure 1: Illustrations of  $\mathcal{R}_g(\mathbf{x})$ . Left:  $g_1 = g_2 = 1$ . Middle:  $g_1 = 1, g_2 = \infty$ . Right: The shape of both the objectives on one of the axes.

### 3. Theoretical Results

In this section we will present our main theoretical results for the relaxations (7) and (8). We will state our results in terms of the vector case (7). However identical results hold for the matrix case when the magnitudes of the vector elements are replaced by the singular values of the matrix.

#### 3.1. What is our relaxation solving?

The goal of our regularization is to adaptively select the appropriate rank/cardinality given the data fidelity. If the true cardinality had been known it would be desirable solve the so called "fixed cardinality" problem

$$\mathbf{o} = \arg \min_{\text{card}(\mathbf{x}) \leq k} \|A\mathbf{x} - \mathbf{b}\|^2. \quad (10)$$

Hence we would like our formulation to determine  $k$  and then to solve (10) exactly. While many regularization methods have been proposed very few of them

output solutions that are "fixed cardinality" minimizers for some  $k$ . As discussed in the previous section most of them add a bias that clearly favors small solutions.

The following theorem gives conditions that ensure that a particular "fixed cardinality" solution is stationary for our relaxation provided that the noise is bounded.

**Theorem 3.1.** *Suppose that  $\mathbf{o}$  solves (10) and  $\tilde{\delta}_k > \sqrt{g_k}$  then  $\mathbf{o}$  is a stationary point (7) if*

$$\|\epsilon\| \leq \frac{\min\{\sqrt{g_{k+1}}, \tilde{\delta}_k\}}{\|A\|}, \quad (11)$$

where  $\epsilon = A\mathbf{o} - \mathbf{b}$  are the residual errors.

The proof of this result and its matrix analog is given in the supplementary material (Appendix B).

There are two essential constraints that ensure that a fixed cardinality solution  $\mathbf{o}$  is stationary. Firstly, since  $\tilde{\delta}_i$  is decreasing and  $\sqrt{g_i}$  is increasing with  $i$  we can view the constraint  $\tilde{\delta}_k > \sqrt{g_k}$  as a threshold  $\sqrt{g_i}$  which must be smaller than any non-zero element. The second constraint essentially states the remaining residual error that is not explained by  $\mathbf{o}$  has to be sufficiently small. As a simple example we mention the noise free case  $\mathbf{b} = A\mathbf{y}$  with  $\text{card}(\mathbf{y}) = k$ . Here  $\epsilon = 0$  and therefore  $\mathbf{x} = \mathbf{y}$  would be a stationary point for all choices of  $g$  where the non-zero elements of  $\mathbf{y}$  are larger than  $\sqrt{g_k}$ .

The above result does not rule out the existence of multiple stationary points. The main results of our paper are dedicated to developing conditions that ensure uniqueness of a stationary point for appropriate choices of  $g$ . In such cases we therefore obtain a method that is able to jointly determine the best  $k$  and supply us with the corresponding "fixed cardinality" solution.

### 3.2. Element Separation and Optimality

The main result of this section will show that a sparse stationary point is under certain conditions unique. The conditions are related to the noise and whether there is a clear truncation of the data or not.

Consider the stationary points of (7). Since  $f^{**}(\mathbf{x}) = \mathcal{R}_g(\tilde{\mathbf{x}}) + \|\tilde{\mathbf{x}}\|^2$  we can write the objective function (1) as  $f^{**}(\mathbf{x}) + h(\mathbf{x})$ , where

$$\begin{aligned} h(\mathbf{x}) &= \|A\mathbf{x} - \mathbf{b}\|^2 - \|\mathbf{x}\|^2 \\ &= \mathbf{x}^T(A^T A - I)\mathbf{x} - 2\mathbf{x}^T A^T \mathbf{b} + \mathbf{b}^T \mathbf{b}, \end{aligned} \quad (12)$$

which has  $\nabla h(\mathbf{x}) = -2\mathbf{z}$ , where  $\mathbf{z} = (I - A^T A)\mathbf{x} + A^T \mathbf{b}$ . A point is stationary if and only if  $2\mathbf{z} = -\nabla h(\mathbf{x}) \in \partial f^{**}(\mathbf{x})$ . Some properties of the solutions to these

equations can be understood by noting that for a fixed  $\mathbf{z}$  the exact same equations are obtained by differentiating the objective

$$\mathcal{R}_g(\tilde{\mathbf{x}}) + \|\mathbf{x} - \mathbf{z}\|^2 = f^{**}(\mathbf{x}) - 2\mathbf{x}^T \mathbf{z} + \|\mathbf{z}\|^2. \quad (13)$$

This expression can be seen as a local approximation (ignoring constants) around  $\mathbf{x}$ , and  $\mathbf{x}$  is stationary in (7) if and only if it is stationary in (13). Furthermore, (13) is the convex envelope of  $G(\text{card}(\mathbf{x})) + \|\mathbf{x} - \mathbf{z}\|^2$ . Therefore the stationary point  $\mathbf{x}$  is the best low cardinality approximation of  $\mathbf{z}$  which is obtained by truncating the elements  $\tilde{z}_i$  at  $\sqrt{g_i}$ .

It is the properties of  $\mathbf{z}$  that decide if there could be other (sparse) stationary points or not. Loosely speaking, our theory relies on the fact that the directional derivative of  $f^{**}$  grows faster than that of  $-h$ . Hence if they are equal at some stationary point they cannot be so again. This is true if the magnitudes  $\tilde{z}_i$  are well separated from their thresholds. If this is not the case a small change in  $\tilde{z}_i$  can cause  $\tilde{x}_i$  to switch from  $\tilde{x}_i = \tilde{z}_i$  to  $\tilde{x}_i = 0$  (or vice versa) which may result in the directional derivative of  $f^{**}$  not growing sufficiently fast. Note that for noise free recovery, that is  $\mathbf{b} = A\mathbf{y}$  for some vector  $\mathbf{y}$ , we have  $\mathbf{z} = (I - A^T A)\mathbf{y} + A^T A\mathbf{y} = \mathbf{y}$ . Hence as long as the non-zero elements of  $\mathbf{y}$  are sufficiently separated from 0 we should be able to guarantee that this is the only sparse stationary point.

We now state the main result:

**Theorem 3.2.** *Suppose that  $\mathbf{x}$  is a stationary point of (7), that is,  $2\mathbf{z} \in \partial f^{**}(\mathbf{x})$  with  $\mathbf{z} = (I - A^T A)\mathbf{x} + A^T \mathbf{b}$ , and the matrix  $A$  fulfills (4). If  $\text{card}(\mathbf{x}) = k$ ,  $\tilde{x}_i \notin (0, \sqrt{g_i})$  and  $\tilde{\mathbf{z}}$  fulfills*

$$\tilde{z}_i \notin \left[ (1 - \delta_r)\sqrt{g_k}, \frac{\sqrt{g_k}}{(1 - \delta_r)} \right] \quad \text{and} \quad \tilde{z}_{k+1} < (1 - 2\delta_r)\tilde{z}_k, \quad (14)$$

then any other stationary point  $\mathbf{x}'$  has  $\text{card}(\mathbf{x}') > r - k$ . If in addition  $k < \frac{r}{2}$  then  $\mathbf{x}$  solves

$$\min_{\text{card}(\mathbf{x}) < \frac{r}{2}} \mathcal{R}_g(\tilde{\mathbf{x}}) + \|A\mathbf{x} - \mathbf{b}\|^2. \quad (15)$$

The proof of this theorem is given in the supplementary material (Appendix C). To gain some more understanding of the conditions (14) we recall that the sequence  $\sqrt{g_i}$  is non-decreasing, while  $\tilde{\mathbf{x}}$  and  $\tilde{\mathbf{z}}$  are non-increasing. Therefore the first condition in (14) ensures that none of the elements  $\tilde{z}_i$  are close to any of their thresholds  $\sqrt{g_i}$ . The second condition is to prevent a change of the permutation  $\pi$ , since this may result in two or more elements in  $\mathbf{x}$  to switch from non-zero to zero and vice versa. This can happen if  $\tilde{z}_k$  and  $\tilde{z}_{k+1}$  are close to each other. In other cases changes in ordering does not cause changes in the support of  $\mathbf{x}$ .

The assumption  $\tilde{x}_i \notin (0, \sqrt{g_i})$  is equivalent to  $f^{**}(\mathbf{x}) = f(\mathbf{x})$ , see supplementary material (Appendix A). In principle there could be stationary points where this assumption is not fulfilled if the regularizer is not strong enough to force such elements to be zero. For example, if the data term is of the form  $\mu \|\mathbf{x} - \mathbf{b}\|^2$ , increasing  $\mu$  will eventually lead to the optimal solution being  $\mathbf{b}$  regardless of the size of its elements. One way to ensure that the regularization is strong enough is to require that  $\|\mathcal{A}\| < 1$ . Then any local minimizer of (7) will have  $f^{**}(\mathbf{x}) = f(\mathbf{x})$  by Theorem 4.7 in [9]. In addition any local minimizer of (7) will be a local minimizer of (1) (but not the other way around) and the global minimizers will coincide.

Before proceeding we also note that since  $f^{**} \leq f$  the stationary point in Theorem 3.2 will also solve

$$\min_{\text{card}(\mathbf{x}) < \frac{r}{2}} G(\text{card}(\mathbf{x})) + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2, \quad (16)$$

if it solves (15), meaning that in some sense it is the best possible sparse solution to the problem. In what follows we will therefore assume that  $r = 2k$ .

We conclude the section by giving results that are sufficient to guarantee the existence of a stationary point fulfilling the constraints of Theorem 3.2 in the presence of noise, that is,  $\mathbf{b} = \mathbf{A}\mathbf{y} + \boldsymbol{\epsilon}$  for some clean vector  $\mathbf{y}$ . The following result shows that as long as the noise level is not too high there will be a stationary point fulfilling the constraints of Theorem 3.2. Moreover, this is true for a whole range of objectives as long as the thresholds  $\sqrt{g_k}$  are not selected too close to the elements in  $\tilde{\mathbf{y}}$ . Note that this result relies on worst case bounds in terms of the noise vector  $\boldsymbol{\epsilon}$ . The proof basically assumes that a single element of the stationary point  $\mathbf{x}$  is affected by the full noise magnitude  $\|\boldsymbol{\epsilon}\|$  rather than evenly distributing the noise among the elements of  $\mathbf{x}$ . This makes the statement weaker than what can be expected in practice with for example Gaussian noise.

**Theorem 3.3.** *Suppose that  $\mathbf{b} = \mathbf{A}\mathbf{y} + \boldsymbol{\epsilon}$ , for some  $\mathbf{y}$  with  $\text{card}(\mathbf{y}) = k$ ,  $\|\mathbf{A}\| < 1$ ,  $\delta_{2k} < \frac{1}{2}$ . If*

$$\tilde{y}_k > \frac{5}{(1 - 2\delta_{2k})\sqrt{1 - \delta_{2k}}} \|\boldsymbol{\epsilon}\|, \quad (17)$$

*then (7) has a stationary point  $\mathbf{x}$ , with  $\text{card}(\mathbf{x}) = k$ , that fulfills (14) for all choices of  $G$  where*

$$\sqrt{g_k} < (1 - \delta_k) \left( \tilde{y}_k - \frac{2\|\boldsymbol{\epsilon}\|}{\sqrt{1 - \delta_{2k}}} \right) \text{ and } \sqrt{g_{k+1}} > \frac{3(1 - \delta_k)}{\sqrt{1 - \delta_{2k}}} \|\boldsymbol{\epsilon}\|. \quad (18)$$

Note that the proof, which is given in the supplementary material (Appendix D), shows that regardless of the choice of  $G$  the stationary point is always the best

cardinality  $k$  approximation of  $\mathbf{y}$  (in a least squares sense). Hence if we know the cardinality beforehand we might as well use the fixed-cardinality relaxation  $g_i = 0$  if  $i \leq k$  and  $g_i = \infty$  if  $i > k$ . In many practical cases the cardinality is not known beforehand but needs to be determined through a suitably selected  $G$  function. The above estimates show that a solution that is close to the original noise free vector  $\mathbf{y}$  (and has the correct support) can often be recovered.

### 3.3. Regularizers with Hard Constraints

The theory presented in the previous section shows uniqueness of sufficiently sparse stationary points, but cannot rule out dense stationary points. The main difficulty in this respect is that the RIP constraint only gives information about low cardinality vectors, and typically there are dense vectors in the nullspace of  $\mathbf{A}$ . As illustrated in Section 2 unbiased separable regularizers don't penalize these vectors sufficiently.

In this section we assume that we know an upper bound  $k_{\max}$  on the cardinality. This means that  $g_i = \infty$  for all  $i \geq k_{\max}$ . The next result shows that relaxations resulting from such regularizers turn out to be strong enough to exclude the existence of high cardinality local minimizers and giving global optimality of a solution fulfilling the assumptions of Theorem 3.2. Note that this prior is by construction non-separable since it counts the number of non-zero element.

**Corollary 3.4.** *Suppose that  $\mathbf{x}$  is a stationary point of (7) that fulfills the assumptions of Theorem 3.2 with  $r = 2k$ . If  $\|\mathbf{A}\| < 1$  and  $g_i = \infty$  for  $i \geq k$  then  $\mathbf{x}$  is the unique local minimizer of (7) and a global minimizer of (1).*

*Proof.* According to Theorem 4.7 of [9] a local minimizer  $\mathbf{x}$  of (7) has  $f^{**}(\mathbf{x}) = f(\mathbf{x})$ . If  $\text{card}(\mathbf{x}) > k$  then  $G(\text{card}(\mathbf{x})) = \infty$  but  $f^{**}(\mathbf{x})$  is finite. Therefore any local minimizer has to have  $\text{card}(\mathbf{x}) \leq k$ . However, according to Theorem 3.2  $\mathbf{x}$  is the only stationary point with  $\text{card}(\mathbf{x}) \leq k$ .  $\square$

With the assumptions of Theorem 3.3 we get the following somewhat stronger result which also ensures existence.

**Corollary 3.5.** *Under the assumptions of Theorem 3.3 the problem (7) has a unique local minimizer which is also a global minimizer of (1).*

## 4. Experiments

In this section we illustrate the behaviour of the proposed penalty using a range of numerical experiments, both synthetic and from real applications. We are in



particular interested in differences between the two relaxations obtained with  $g_i = \mu$  for all  $i$  versus

$$\bar{g}_i = \begin{cases} \mu & i \leq k_{\max} \\ \infty & i > k_{\max}. \end{cases}$$

We denote the functions obtained with these choices  $G_\mu(\text{card})$  and  $\bar{G}(\text{card})$  respectively and their relaxations  $\mathcal{R}_\mu$  and  $\mathcal{R}_{\bar{g}}$  respectively. Both of these attempt to estimate an unknown cardinality based on a trade-off between data fidelity and sparsity. However the second option  $\mathcal{R}_{\bar{g}}$  should be more robust to local minima with high cardinality as outlined in our theory.

## 4.1. Synthetic data

### 4.1.1 Robustness - random matrices

To test the robustness with respect to noise we generated 10 different problem instances - i.e. different  $100 \times 200$  Gaussian random matrices  $A$  (with normalized columns), sparse real ground truths  $\mathbf{x}_0$  and noise  $\epsilon$  - for each noise level  $\|\epsilon\|/\|\mathbf{b}\| \in \{0.025i : i = 0, \dots, 10\}$  and averaged the output distances  $\|\tilde{\mathbf{x}} - \mathbf{x}_0\|/\|\mathbf{x}_0\|$ , with  $\tilde{\mathbf{x}}$  approximated solution computed by the minimizing algorithm. Each  $\mathbf{x}_0$  was chosen with random cardinality between 10 and 18 and with the property that  $\min_i |\mathbf{x}_{0,i}| > 2\sqrt{2}$ . To approximately recover  $\mathbf{x}_0$  we use the formulation  $g_i = 2$  for  $i \leq k_{\max} = 20$  and  $g_i = \infty$  otherwise, and refer to it as  $\mathcal{R}_{\bar{g}}$  in the results. We used  $\mathbf{0}$  as starting point and the Forward-Backward Splitting as minimizing algorithm.

For comparison we also test the Least Absolute Shrinkage and Selection Operator (LASSO),  $\ell^p$  with  $p = 1/2, 2/3$ , and the Smoothly Clipped Absolute Deviation (SCAD) which are popular approaches from the literature. To avoid issues with parameter selection and achieve the best possible performance for these competing methods their parameters were picked using a line search *for each problem instance*. So, for each problem instance  $A$ ,  $\epsilon$  and  $\mathbf{x}_0$ , we computed  $\|\tilde{\mathbf{x}} - \mathbf{x}_0\|/\|\mathbf{x}_0\|$  for several choices of the involved parameters and stored the best outcome. The result is shown in the top left graph in Figure 2. Here  $\mathcal{R}_g$  outperforms the other relaxations consistently giving the best fit to the ground truth data. The behaviour that all methods approach the ground truth solution when the noise decreases is due to the fact that the parameter is exactly tuned to the noise level for each problem. We emphasize that in real applications such a strategy is normally not feasible when the noise level is unknown.

A more realistic scenario is to use the same parameter setting for all noise levels. Therefore we also performed a second batch of experiments where one single parameter for each competitor method was used. The

parameters chosen minimize the average error through all the noise levels and all trials. The outcome (second row of Figure 2) highlights the benefits of the noise-invariance of the bias-free methods ( $\mathcal{R}_g$  and SCAD). These can handle varying noise levels with a single parameter setting.

### 4.1.2 Robustness - concatenation between Fourier transform and identity

For a matrix  $A \in \mathbb{C}^{m \times N}$ , the quantity

$$\mu(A) = \max_{1 \leq i \neq j \leq N} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$$

is known as *mutual coherence* of  $A$ . In compressive sensing a small mutual coherence is desirable because it controls all the restricted isometry constants from above (cfr. Prop. 6.2 in [18]); the matrix  $A = [F|I]$  with  $F$  being the  $1D$  Fourier transform matrix and  $I$  the identity matrix, is known to have mutual coherence  $= 1/\sqrt{m}$  and it is often used in compressive sensing algorithms or techniques benchmarking.

We ran a similar set of tests as in Section 4.1.1 using  $A$  of dimensions  $100 \times 200$  instead of random matrices: we generated 10 different problem instances - i.e. different sparse real ground truths  $\mathbf{x}_0$  and noise  $\epsilon$  - for each noise level  $\|\epsilon\|/\|\mathbf{b}\| \in \{0.025i : i = 0, \dots, 10\}$  and averaged the output distances  $\|\tilde{\mathbf{x}} - \mathbf{x}_0\|/\|\mathbf{x}_0\|$ , with  $\tilde{\mathbf{x}}$  approximated solution computed by the minimizing algorithm. Each  $\mathbf{x}_0$  was chosen with cardinality 10 and with the property that  $\min_i |\mathbf{x}_{0,i}| > 2\sqrt{2}$ . In our formulation we selected  $g_i = 2$  for  $i = 1, \dots, k_{\max} = 20$ ; the competitor methods are the same as in Section 4.1.1 and their parameters choices were again made via the same technique(s). The outcome mostly mirrors that in Section 4.1.1, see the second columns of Figure 2.

### 4.1.3 Sparsity

In a second batch of experiments we studied the cardinality of the retrieved approximation when a starting point not too far from the ground truth is employed. For a fixed noise level  $\|\epsilon\|/\|\mathbf{b}\| = 0.15$  we fixed a triplet  $A$  (Gaussian, with normalized columns),  $\epsilon$  and  $\mathbf{x}_0$  with  $\text{card}(\mathbf{x}_0) = 10$  and again  $\min_i |\mathbf{x}_{0,i}| > 2\sqrt{2}$ , and we generated 250 different random (with uniform distribution) starting points  $\mathbf{x}_S$  with  $0.2\|\mathbf{x}_0\| \leq \|\mathbf{x}_S\| \leq 3\|\mathbf{x}_0\|$ . In Table 1 we display mean and standard deviation of  $S_m(\tilde{\mathbf{x}}, \mathbf{x}_0) = \text{card}(\text{supp}(\tilde{\mathbf{x}}) \Delta \text{supp}(\mathbf{x}_0))^2$  and normalized distance to ground truth  $\|\tilde{\mathbf{x}} - \mathbf{x}_0\|/\|\mathbf{x}_0\|$  in the scenario  $75 \times 200$ . The parameter choice for the competing methods was made with a line

<sup>2</sup> $\Delta$  is the set-theoretic symmetric difference.

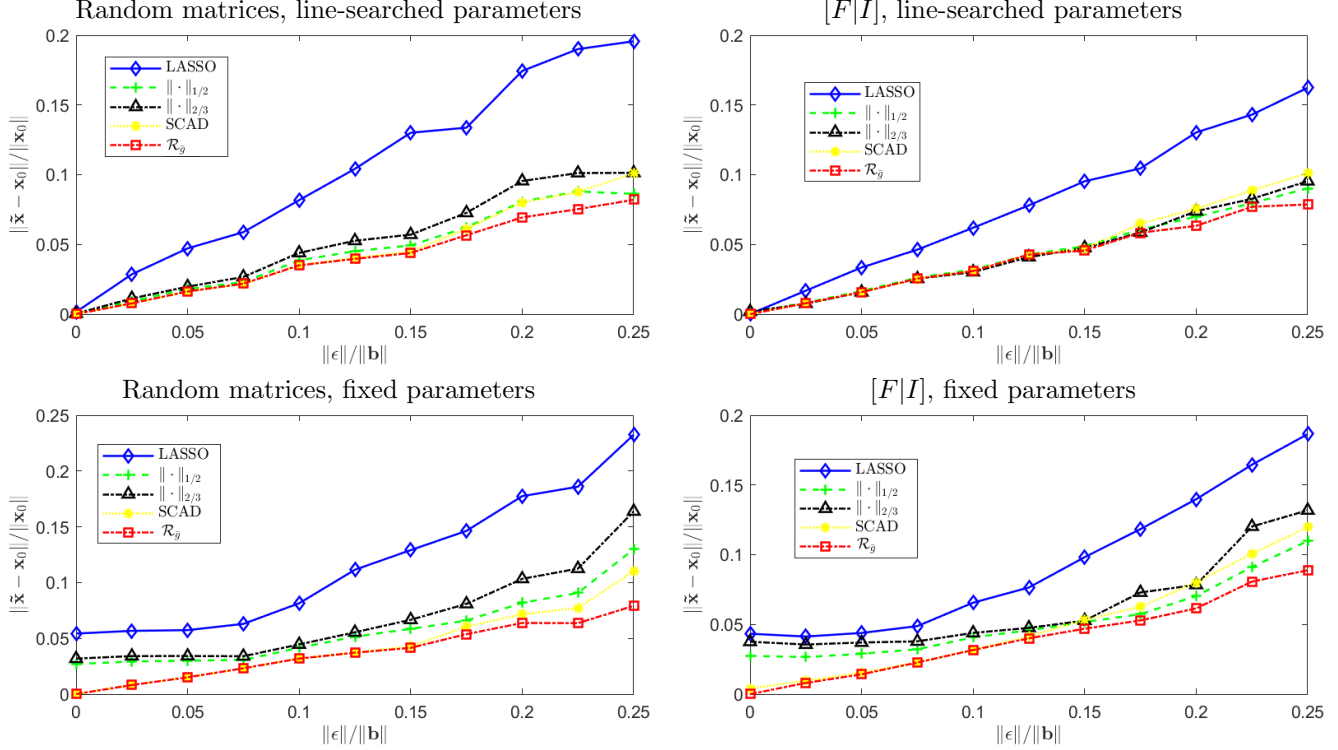


Figure 2: Normalized error vs normalized noise level. The first column experiments with random matrices, see Section 4.1.1; the second, concatenation of Fourier matrix with identity, see Section 4.1.2.

search for each single problem instance (as in the previous section). To achieve a the correct support as well as a good fit to the ground truth solution we selected the parameter that minimized the quantity  $0.8 \text{card}(\text{supp}(\tilde{\mathbf{x}}) \Delta \text{supp}(\mathbf{x}_0)) / \|\mathbf{x}_0\|_0 + 0.2 \|\tilde{\mathbf{x}} - \mathbf{x}_0\| / \|\mathbf{x}_0\|$ . The proposed regularizer (indicated in the legend as “ $\mathcal{R}_{\bar{g}}$ ”) displays a solid behaviour.

Table 1: Random starting point,  $75 \times 200$  scenario

	$\ \tilde{\mathbf{x}} - \mathbf{x}_0\  / \ \mathbf{x}_0\ $		$S_m(\tilde{\mathbf{x}}, \mathbf{x}_0)$	
	Mean	St. dev.	Mean	St. dev.
$\mathcal{R}_{\bar{g}}$	0.0768	$1.96 \cdot 10^{-7}$	0	0
$\mathcal{R}_{\mu}$	0.0768	$2.20 \cdot 10^{-6}$	0	0
SCAD	0.0898	$2.03 \cdot 10^{-7}$	0	0
LASSO	0.7290	$1.09 \cdot 10^{-7}$	2	0
$\ell^{1/2}$	0.1465	0.1626	0.4840	1.6915
$\ell^{2/3}$	0.1147	0.0052	0	0

#### 4.1.4 Local minima suppression

In [12] was numerically shown that the algorithm minimizing the functional  $\mathcal{R}_{\mu}(\mathbf{x}) + \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2$  tends to get stuck in high cardinality local minima when the

Table 2:  $\mathbf{A}^{\dagger}\mathbf{b}$  as starting point,  $100 \times 200$  scenario

	$\ \tilde{\mathbf{x}} - \mathbf{x}_0\  / \ \mathbf{x}_0\ $		$S_m(\tilde{\mathbf{x}}, \mathbf{x}_0)$	
	Mean	St. dev.	Mean	St. dev.
$\mathcal{R}_{\bar{g}}$	0.0643	0.1079	0.1080	0.3273
$\mathcal{R}_{\mu}$	0.0636	0.1343	0.2360	2.9772
SCAD	0.0885	0.0582	0.1240	0.3758
LASSO	0.3863	0.1490	1.2800	1.5734
$\ell^{1/2}$	0.0856	0.1449	0.4560	2.4658
$\ell^{2/3}$	0.0650	0.0157	0	0

Table 3: Local minima counter

	$\mathcal{R}_{\bar{g}}$	$\mathcal{R}_{\mu}$	$\ \cdot\ _{1/2}$	$\ \cdot\ _{2/3}$
Loc. min. detected	0	100	1	0

least square solution is used as starting point. The present paper can be also seen as an attempt to overcome that issue and in this section we numerically confirm that those high cardinality local minima seem to be suppressed when our new penalty is employed. We again generated 250 different problem instances with the same specs as in Section 4.1.3 and used a least




Figure 3: Rank of  $X^\#$  (x-axis) versus data fit  $\|RX - M\|_F$  (y-axis) for the MOCAP sequences, drink, pickup, stretch and yoga, used in [14]. Blue -  $\mathcal{R}_g$ , red -  $\|\cdot\|_*$ , yellow - SCAD, purple - Schatten-2/3 and green - Schatten-1/2.

square solution to the linear system  $A\mathbf{x} = \mathbf{b}$  as starting point for the algorithm. Note that since  $A$  has a nullspace there are in general multiple least squares solutions. In Table we 2 used  $A^\dagger\mathbf{b}$ . The results show that while  $A^\dagger\mathbf{b}$  seems like a sensible starting point it often gives sub-optimal results. This is in particular true for the bias free regularizer  $\mathcal{R}_\mu$  that has difficulty recovering from a high cardinality starting point. The proposed  $\mathcal{R}_{\bar{g}}$  is in general much less affected than the other bias free methods. We also remark that strictly speaking deviations from the ground truth  $x_0$  may not be a result of local minima since  $x_0$  is not a minimizer.

The point  $A^\dagger\mathbf{b}$  is an intuitive initialization and indeed it is a bit surprising that it produces local minima. A less intuitive choice is what is described in Section 2, that is, points of type  $A^\dagger\mathbf{b} + \mathbf{x}_k$  where  $\mathbf{x}_k \in \ker(A)$  is dense such that  $\min_i |(A^\dagger\mathbf{b} + \mathbf{x}_k)_i| > 2\sqrt{2}$ , which we consider in Table 3. These are still least square solutions and they are located in the region where the penalty  $\mathcal{R}_\mu(\mathbf{x})$  is constant; thus they are local minima for the functional  $\mathcal{R}_\mu(\mathbf{x}) + \|A\mathbf{x} - \mathbf{b}\|^2$ . For a  $100 \times 200$  random matrix  $A$  we generated 100 linearly independent  $\mathbf{x}_k$  and tested whether the points  $A^\dagger\mathbf{b} + \mathbf{x}_k$  are local minima or not for the functionals  $\mathcal{R}_\mu(\mathbf{x}) + \|A\mathbf{x} - \mathbf{b}\|^2$ , (7) and  $\|\mathbf{x}\|_p + \|A\mathbf{x} - \mathbf{b}\|^2$  (for  $p = 1/2, 2/3$ ); we here picked  $k_{\max} = 16 > \text{card}(\mathbf{x}_0) \in \{9, 10, \dots, 14\}$  chosen randomly. Table 3 displays the results of our experiment: it shows that all those points are (as expected) local minima for  $\mathcal{R}_\mu(\mathbf{x}) + \|A\mathbf{x} - \mathbf{b}\|^2$  but not for (7) and motivates one more time the constructions in the present manuscript.

## 4.2. Rank Regularization for NRSfM

We conclude our experiments by considering an application of the matrix version of our framework. Non-rigid structure from motion (NRSfM) is a classical computer vision problem where object dynamics is modeled using a rank constraint. In this section we follow [14] and extract a deforming model from point tracks obtained with a moving camera. Under the linear deformation assumption [5, 14] the deforming 3D point

cloud can be represented using a low rank matrix  $S$  where row  $i$  contains  $x$ -,  $y$ - and  $z$ - coordinates of the point cloud when image  $i$  was captured. To recover  $S$  we solve

$$\min_S R_g(S) + \|RS^\# - M\|^2. \quad (19)$$

Here  $R$  is a matrix containing camera rotations and  $S^\#$  is a matrix where the elements of  $S$  have been re-ordered, see [14] for detailed definitions.

In Figure 3 we use data from [14] to test our regularizer  $\mathcal{R}_g$ , with  $g_k = 0$  if  $k \leq k_{\max}$  and  $\infty$  otherwise, against the nuclear norm, SCAD and the Schatten-norms. Only  $\mathcal{R}_g$  can directly penalize the rank. The other competing methods have different parameters that needs to be tuned to indirectly obtain a certain rank. Not that while a small parameter change may not change the rank it can still change the solution since the bias is affected. For a fair comparison we therefore sample parameters over a whole range of values to see what data fit can we achieve with settings that give particular rank. Note that since the operator  $S \mapsto RS^\#$  has low rank matrices in its null-space it therefore does not fulfill RIP, therefore our method could have local minima. For ranks between 2 and 6 Figure 3 shows error bars covering the best and the worst data fit for each method. The bias is most clearly visible for the lowest rank (2) where the methods have to suppress more noise. Our method (blue) consistently gives the lowest data fit for each rank.

## 5. Conclusions

In this paper we have presented and analysed a general framework for sparsity and rank regularization of linear least squares problems. Our regularizers are bias free and non-separable which admits increased modeling power compared to standard separable version. Our theoretical analysis shows that under the RIP constraint stationary points are often unique even though our framework is non-convex. Our empirical results further demonstrate that we outperform competing methods in terms of accuracy and robustness.



## References

- [1] Andreas Argyriou, Rina Foygel, and Nathan Srebro. Sparse prediction with the k-support norm. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1457–1465. Curran Associates, Inc., 2012. [2](#)
- [2] Ronen Basri, David Jacobs, and Ira Kemelmacher. Photometric stereo with general, unknown lighting. *International Journal of Computer Vision*, 72(3):239–257, May 2007. [1](#)
- [3] Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *Applied and computational harmonic analysis*, 27(3):265–274, 2009. [2](#)
- [4] Kristian Bredies, Dirk A. Lorenz, and Stefan Reiterer. Minimization of non-smooth, non-convex functionals by iterative thresholding. *Journal of Optimization Theory and Applications*, 165(1):78–112, 2015. [2](#)
- [5] C. Bregler, A. Hertzmann, and H. Biermann. Recovering non-rigid 3d shape from image streams. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2000. [1](#), [8](#)
- [6] Emmanuel J. Candès, Xiaodong Li, Yi Ma, and John Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, 2011. [2](#)
- [7] Emmanuel J. Candès and Benjamin Recht. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics*, 9(6):717–772, 2009. [2](#)
- [8] Emmanuel J. Candès and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE transactions on information theory*, 52(12):5406–5425, 2006. [1](#)
- [9] Marcus Carlsson. On convex envelopes and regularization of non-convex functionals without moving global minima. *Journal of Optimization Theory and Applications*, to appear, 2019. [2](#), [5](#)
- [10] Marcus Carlsson, Daniele Gerosa, and Carl Olsson. Bias reduction in compressed sensing, 2018. [3](#)
- [11] Marcus Carlsson, Daniele Gerosa, and Carl Olsson. An un-biased approach to low rank recovery. *arXiv preprint arXiv:1909.13363*, 2019. [15](#)
- [12] Marcus Carlsson, Daniele Gerosa, and Carl Olsson. An unbiased approach to compressed sensing. *Inverse Problems*, 36 115014, 2020. [7](#)
- [13] Rick Chartrand. Exact reconstruction of sparse signals via nonconvex minimization. *IEEE Signal Processing Letters*, 14(10):707–710, 2007. [2](#)
- [14] Yuchao Dai, Hongdong Li, and Mingyi He. A simple prior-free method for non-rigid structure-from-motion factorization. *International Journal of Computer Vision*, 107(2):101–122, 2014. [8](#)
- [15] David L. Donoho and Michael Elad. Optimally sparse representation in general (non-orthogonal) dictionaries via  $l_1$ -minimization. In *PROC. NATL ACAD. SCI. USA 100 2197–202*, 2002. [1](#)
- [16] Anders Eriksson, Trung Thanh Pham, Tat-Jun Chin, and Ian Reid. The k-support norm and convex envelopes of cardinality and rank. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2015. [2](#)
- [17] Jianqing Fan and Runze Li. Variable selection via nonconcave penalized likelihood and its oracle properties. *Journal of the American Statistical Association*, 96(456):1348–1360, 2001. [2](#)
- [18] Simon Foucart and Holger Rauhut. A mathematical introduction to compressive sensing. 2013. [6](#)
- [19] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. Dense variational reconstruction of non-rigid surfaces from monocular video. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013. [1](#)
- [20] Ravi Garg, Anastasios Roussos, and Lourdes Agapito. A variational approach to video registration with subspace constraints. *International Journal of Computer Vision*, 104(3):286–314, 2013. [1](#)
- [21] N. Gillis and F. Glinier. Low-rank matrix approximation with weights or missing data is np-hard. *SIAM Journal on Matrix Analysis and Applications*, 32(4), 2011. [1](#)
- [22] Christian Grussler, Anders Rantzer, and Pontus Giselsson. Low-rank optimization with convex constraints. *IEEE Transactions on Automatic Control*, 63(11):4000–4007, 2018. [2](#)
- [23] Shuhang Gu, Qi Xie, Deyu Meng, Wangmeng Zuo, Xiangchu Feng, and Lei Zhang. Weighted nuclear norm minimization and its applications to low level vision. *International Journal of Computer Vision*, 121, 07 2016. [2](#)
- [24] Yao Hu, Debing Zhang, Jieping Ye, Xuelong Li, and Xiaofei He. Fast and accurate matrix completion via truncated nuclear norm regularization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(9):2117–2130, 2013. [2](#)
- [25] Jose Pedro Iglesias, Carl Olsson, and Marcus Valtonen Örnåhag. Accurate optimization of weighted nuclear norm for non-rigid structure from motion, 2020. [2](#)
- [26] Suryansh Kumar. A simple prior-free method for non-rigid structure-from-motion factorization : Revisited. *CoRR*, abs/1902.10274, 2019. [2](#)
- [27] Viktor Larsson and Carl Olsson. Convex low rank approximation. *International Journal of Computer Vision*, 120(2):194–214, 2016. [2](#)
- [28] Po-Ling Loh and Martin J. Wainwright. Regularized m-estimators with nonconvexity: Statistical and algorithmic theory for local optima. In *Advances in Neural*

- Information Processing Systems*, pages 476–484, 2013. [2](#)
- [29] Po-Ling Loh and Martin J. Wainwright. Support recovery without incoherence: A case for nonconvex regularization. *arXiv preprint*, arXiv:1412.5632, 2014. [2](#)
- [30] Po-Ling Loh and Martin J Wainwright. Support recovery without incoherence: A case for nonconvex regularization. *The Annals of Statistics*, 45(6):2455–2482, 2017. [2](#)
- [31] Rahul Mazumder, Jerome H. Friedman, and Trevor Hastie. Sparsenet: Coordinate descent with nonconvex penalties. *Journal of the American Statistical Association*, 106(495):1125–1138, 2011. [2](#)
- [32] Andrew M. McDonald, Massimiliano Pontil, and Dimitris Stamos. New perspectives on k-support and cluster norms. *J. Mach. Learn. Res.*, 17(1):5376–5413, January 2016. [2](#)
- [33] Karthik Mohan and Maryam Fazel. Iterative reweighted least squares for matrix rank minimization. In *Annual Allerton Conference on Communication, Control, and Computing*, pages 653–661, 2010. [2](#)
- [34] Balas K. Natarajan. Sparse approximate solutions to linear systems. *SIAM journal on computing*, 24(2):227–234, 1995. [1](#)
- [35] Tae-Hyun Oh, Yu-Wing Tai, Jean-Charles Bazin, Hyeonwoo Kim, and In S. Kweon. Partial sum minimization of singular values in robust pca: Algorithm and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(4):744–758, 2016. [2](#)
- [36] Carl Olsson, Marcus Carlsson, Fredrik Andersson, and Viktor Larsson. Non-convex rank/sparsity regularization and local minima. *Proceedings of the International Conference on Computer Vision*, 2017. [2](#), [3](#)
- [37] Carl Olsson, Marcus Carlsson, and Erik Bylow. A non-convex relaxation for fixed-rank approximation. In *2017 IEEE International Conference on Computer Vision Workshops (ICCVW)*, pages 1809–1817, Oct 2017. [2](#)
- [38] Carl Olsson, Anders Eriksson, and Richard Hartley. Outlier removal using duality. In *IEEE Int. Conference on Computer Vision and Pattern Recognition*, pages 1450–1457, 2010. [1](#)
- [39] Samet Oymak, Amin Jalali, Maryam Fazel, Yonina C. Eldar, and Babak Hassibi. Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Transactions on Information Theory*, 61(5):2886–2908, 2015. [2](#)
- [40] Samet Oymak, Karthik Mohan, Maryam Fazel, and Babak Hassibi. A simplified approach to recovery conditions for low rank matrices. In *IEEE International Symposium on Information Theory Proceedings (ISIT)*, pages 2318–2322, 2011. [2](#)
- [41] Zheng Pan and Changshui Zhang. Relaxed sparse eigenvalue conditions for sparse estimation via nonconvex regularized regression. *Pattern Recognition*, 48(1):231–243, 2015. [2](#)
- [42] Benjamin Recht, Maryam Fazel, and Pablo A. Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.*, 52(3):471–501, August 2010. [2](#)
- [43] Carlo Tomasi and Takeo Kanade. Shape and motion from image streams under orthography: A factorization method. *International Journal of Computer Vision*, 9(2):137–154, 1992. [1](#)
- [44] Joel A. Tropp. Just relax: Convex programming methods for identifying sparse signals in noise. *IEEE transactions on information theory*, 52(3):1030–1051, 2006. [1](#)
- [45] Joel A. Tropp. Convex recovery of a structured signal from independent random linear measurements. In *Sampling Theory, a Renaissance*, pages 67–101. 2015. [1](#)
- [46] John Wright, Allen Y. Yang, Arvind Ganesh, S. Shankar Sastry, and Yi Ma. Robust face recognition via sparse representation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 31(2):210–227, February 2009. [1](#)
- [47] Jingyu Yan and Marc Pollefeys. A factorization-based approach for articulated nonrigid shape, motion and kinematic chain recovery from video. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(5):865–877, 2008. [1](#)
- [48] Cun-Hui Zhang. Nearly unbiased variable selection under minimax concave penalty. *The Annals of Statistics*, 38(2):894–942, 2010. [2](#)
- [49] Cun-Hui Zhang and Tong Zhang. A general theory of concave regularization for high-dimensional sparse estimation problems. *Statistical Science*, pages 576–593, 2012. [2](#)
- [50] Hui Zou and Runze Li. One-step sparse estimates in nonconcave penalized likelihood models. *Annals of statistics*, 36(4):1509, 2008. [2](#)